Exploring a Network Model of the Mouse Primary Visual Cortex

Lena Myklebust

Thesis submitted for the degree of
Master in Physics
60 credits

Department of Physics
Faculty of mathematics and natural sciences

UNIVERSITY OF OSLO

May 2020

Exploring a Network Model of the Mouse Primary Visual Cortex

Lena Myklebust

Exploring a Network Model of the Mouse Primary Visual Cortex

http://www.duo.uio.no/

Printed: Reprosentralen, University of Oslo

Abstract

A current challenge within neuroscience is to model the dynamics of biological networks. This thesis explores a model of the mouse primary visual cortex (V1) and lateral geniculate nucleus (LGN), developed by the Allen Institute for Brain Science. Spatial and temporal tuning properties of model cells are analysed. Of main focus is two phenomena which have been observed in mice LGN and V1. The first, surround suppression, refers to the neural response to a visual stimulus being inhibited by a surrounding stimulus. The second is band-pass tuning in response to stimuli of various spatial frequencies.

The original Allen model fails to exhibit both surround suppression and band-pass tuning. It is therefore explored how the model could be improved, specifically by modifying the LGN. To represent the center-surround organisation of receptive fields found in mice, the Gaussian spatial filters of the LGN are replaced with Difference-of-Gaussians (DoG). As a result, the modified model reproduces band-pass spatial frequency tuning for grating stimuli and surround suppression for uniform stimuli. However, no clear suppression is shown in response to patches of grating stimuli, inconsistent with both previous findings in mice V1 and the LGN and V1 of other mammals.

Surround suppression is thought to incorporate both linear and nonlinear mechanisms, where only the linear response can be accounted for by the DoG function. How nonlinear suppression arises in the cortex is not yet fully understood, though one potential source is input from LGN. Previous studies have suggested that responses in the visual pathway are effectively normalised through neural inhibition. To model this, the Allen LGN is modified so that firing rates are normalised with respect to overall activity. Suppression is thus reproduced in both LGN and V1 for patch grating stimuli. However, the method has limitations when it comes to accounting for all aspects of V1 suppression, such as time delay and dependence on stimuli orientation. Another potential source of V1 suppression is network effects from within cortex. Recurrent connections in the Allen V1, however, seem to have a limited effect on size tuning. Future studies might therefore want to explore whether changes to these connections could generate suppression.

For temporal frequency tuning of the Allen model, the results are closer to experiments than for spatial tuning. However, the preferred frequencies in the Allen V1 are somewhat higher than what was expected based on previous experiments on mice. There is also overall less variation in tuning properties between model cells in V1 than what has been observed experimentally, both regarding temporal and spatial tuning. A wider range of tuning properties is displayed by the LGN, but the variation seems to average out in V1. In the future, a greater variation in V1 frequency tuning may be obtained by adjusting the cortical and thalamocortical connections according to their preferred frequencies.

Acknowledgements

I would first like to thank my supervisor, Gaute Einevoll, for introducing me to computational neuroscience, and for your help and motivation during my master studies. Thank you for always being available, and for taking the time to answer my many questions during our meetings.

I would like to express my thanks to Torbjørn Vefferstad Ness, Alexander Johannes Stasik and Espen Hagen for your time spent helping with the BMTK and setting up the Allen model. I also want to thank Atle Eskeland Rimehaug, for sharing your knowledge about the model. To Camilla Marie Baastad, for working together in exploring the field of neuroscience, and for our breaks during long days in the maths building and at CINPLA.

To Julie Uv, for being an inspiration and making my days brighter. Thank you for taking a genuine interest in my work, and for our academic discussions.

A special thank you to Rosa Brunskill, for making me laugh and for being someone I can share everything with. Your insights, support and unique friendship is truly appreciated.

I also want to thank my parents and my friends from home for your company and your input on my work. In particular Elin Dypvik Sødahl and Bera Riiser Moe for your support and our study sessions at the university.

My deepest appreciation goes to Robin Pearson, for being there for me every day, even from across the pond. Your good spirits and continuous encouragement has been invaluable over the last years.

Contents

1 Introduction **1**
- 1.1 Motivation . . . 1
- 1.2 The Primary Visual Cortex Model . . . 2
- 1.3 Thesis Structure . . . 2

2 Neuroscience Background **4**
- 2.1 The Mouse Visual Pathway . . . 4
- 2.2 The Neuron - Structure and Properties . . . 7
  - 2.2.1 The nerve cell . . . 7
  - 2.2.2 Receptive fields and center-surround organisation . . . 8

3 Computational Neuroscience **10**
- 3.1 Mechanistic Neuron Modelling . . . 10
  - 3.1.1 The Hodgkin-Huxley model . . . 10
  - 3.1.2 The Generalised Leaky Integrate-and-Fire model . . . 11
  - 3.1.3 Spatially extended models . . . 13
- 3.2 Statistical Neuron Modelling . . . 13
  - 3.2.1 The Linear-Nonlinear-Poisson model . . . 13
- 3.3 Modelling Networks . . . 15

4 Spatial and Temporal Neuronal Response Properties **17**
- 4.1 Size Tuning . . . 17
  - 4.1.1 Linear and nonlinear surround suppression . . . 18
  - 4.1.2 Suppression index . . . 20
- 4.2 Spatial Frequency Tuning . . . 21
- 4.3 Temporal Frequency Tuning . . . 23

5 Modelling Visual Receptive Fields **24**
- 5.1 Separable Receptive Fields . . . 24
- 5.2 The Difference-of-Gaussians Filter . . . 25
- 5.3 Responses of Spatial Filters . . . 26
  - 5.3.1 Response to full-field drifting gratings . . . 26
  - 5.3.2 Response to flashing spots . . . 27
  - 5.3.3 Response to patch gratings . . . 28
- 5.4 The Gabor Filter . . . 30

6 The Allen Model **31**
- 6.1 The Allen Network Model . . . 31
- 6.2 The Allen LGN Model . . . 34

7 Experimental Findings **36**
7.1 Spatial Tuning in the Mouse . . . 36
7.1.1 LGN . . . 36
7.1.2 V1 . . . 40
7.1.3 Effect of anaesthesia . . . 44
7.2 Temporal Tuning in the Mouse . . . 45
7.2.1 LGN . . . 45
7.2.2 V1 . . . 46
7.2.3 Effect of anaesthesia . . . 47

8 Simulation Method **49**
8.1 The Versions of the Allen Model . . . 49
8.1.1 The M1 model . . . 50
8.1.2 The M1$*$ models . . . 51
8.1.3 The M1_{norec} model . . . 51
8.1.4 The M2 model . . . 51
8.2 Running the Models . . . 52
8.3 Visual Stimuli . . . 53
8.4 Calculating Firing Rates . . . 59
8.4.1 Time average vs average modulation . . . 59
8.4.2 LGN firing rates . . . 60
8.4.3 V1 firing rates . . . 62
8.5 Spatial Frequency Tuning . . . 63
8.6 Temporal Frequency Tuning . . . 64
8.7 Size Tuning . . . 64

9 Size and Spatial Frequency Tuning of the Allen Model **66**
9.1 The Original Allen Model . . . 67
9.1.1 Spatial frequency tuning - full-field gratings . . . 67
9.1.2 Size tuning - flashing spots . . . 71
9.1.3 Size tuning - patch gratings . . . 75
9.2 The M1 Allen Model . . . 77
9.2.1 Motivation for modifying the model . . . 77
9.2.2 The DoG filter . . . 77
9.2.3 Analytical parameter tuning: peak LGN frequencies . . . 78
9.2.4 Spatial frequency tuning - full-field gratings . . . 79
9.2.5 Size tuning - flashing spots . . . 83
9.2.6 Size tuning - patch gratings . . . 87
9.3 DoG Filters with Various Parameters . . . 89
9.4 Contributions to Size Tuning in V1 . . . 90
9.4.1 Effect of LGN input . . . 90
9.4.2 Effect of recurrent connections in V1 . . . 90
9.5 Modelling Nonlinear Surround Suppression . . . 92
9.5.1 Divisive normalisation in LGN . . . 93

10 Temporal Frequency Tuning of the Allen Model **97**
10.1 LGN . . . 98
10.2 V1 . . . 99

11 Summary, Conclusion and Future Work **103**
11.1 Size and Spatial Frequency Tuning . . . 103
11.1.1 Original model . . . 103
11.1.2 M1 models . . . 104
11.1.3 Origin of tuning properties in V1 . . . 105
11.1.4 Modelling nonlinear suppression . . . 106
11.2 Temporal Frequency Tuning . . . 108
11.3 Conclusion . . . 109
11.4 Further Research . . . 109

Appendix A Derivations of Filter Responses **122**

Appendix B Figures **124**
B.1 Original Allen Model . . . 124
B.1.1 Spatial frequency tuning - full-field gratings . . . 124
B.1.2 Size tuning - flashing spots . . . 129
B.1.3 Size tuning - patch gratings . . . 130
B.2 M1 Allen Model . . . 133
B.2.1 Spatial frequency tuning - full-field gratings . . . 133
B.2.2 Size tuning - flashing spots . . . 137
B.2.3 Size tuning - patch gratings . . . 140
B.3 DoG Filters with Various Parameters . . . 151
B.4 V1 Without Recurrent Connection . . . 153
B.5 Modelling nonlinear suppression . . . 156
B.6 Temporal frequency tuning . . . 157

Appendix C Code **162**
C.1 Example Code . . . 162

Chapter 1

Introduction

1.1 Motivation

Understanding the underlying mechanisms of brain activity is a primary focus in neuroscience [1]. Most of our current knowledge about the brain is obtained through experimental cell recordings [2]. Descriptive and explanatory models have later been developed based on observations of neuronal structures and activity. In 1963, Hodgkin and Huxley received the Nobel Prize for their discoveries and ideas related to neuronal signalling [3]. In 1952, based on experimental evidence using a single squid axon, they proposed an explanation for how the neuron action potential is generated [4–8]. They also introduced a mathematical framework for describing these mechanisms [4–8]. Their findings represent a milestone in neuroscience and the starting point of detailed mathematical neuron models [9].

Simpler, more computationally efficient models have also been developed, such as the Integrate-and-Fire neuron, which makes up the neuron models used in this thesis. Although simple, the Integrate-and-Fire neuron is a good model for gaining insight into the neuronal dynamics of the brain. Predictions from a single Integrate-and-Fire neuron have been found to agree well with experimental recordings [9]. In addition, networks of Integrate-and-Fire models are fairly good at reproducing qualitative aspects captured by more biologically detailed models [10–12].

With recent advances in computational power and speed, simulations of large, complex neuronal networks have been made feasible [13]. Being able to model processes in the brain is arguably necessary in order to understand the cortical functions behind experimental data [12]. Additionally, if a model is representative of the real system, it could be used to as an alternative to experiments [13]. One of the research institutes aiming to develop such models is the Allen Institute for Brain Science[1].

A large amount of empirical data is essential when aiming to create a realistic brain model [14]. In 2019, after performing several electrical recordings and integrating data from years of empirical neuroscience, the Allen Institute for Brain Science released a model of the primary visual cortex in the mouse [12]. The model is a part of the institute's work towards understanding the human brain, and improving health by gaining insight into illnesses, disorders and possible cures [12]. The model is still under development, and several aspects of it remain to be explored [12]. Using experimental data as a reference, its performance may be evaluated, regarding various stimuli and

[1]Further information about the Allen Institute and what they do is available on their web page: `https://alleninstitute.org/what-we-do/brain-science/`

phenomena. This thesis investigates the model's response to stimuli of different sizes, spatial frequencies and temporal frequencies. In particular, two properties which have been observed in the lateral geniculate nucleus and primary visual cortex of real mice are considered. The first, surround suppression, is defined as the neural activity in response to a visual stimulus being inhibited by a surrounding stimulus [15]. The second is band-pass tuning of responses as a function of stimulus spatial frequency. After considering the original model, the effects of modifying the LGN on these properties is explored.

1.2 The Primary Visual Cortex Model

The model explored and tested in this thesis represents the mouse primary visual cortex (V1) as a network of simplified neurons [12]. From here on, this will be referred to as the Allen model. The model is available at two granularities: One where the neurons are represented with spatially extended structures, and one where each cell is reduced to a point [12]. The latter, made up of Generalised Integrate-and-Fire neurons [12], is explored in this thesis. The cells in the network are developed based on extensive experimental research, and are divided into 17 different neuron classes [12]. In addition to the V1 network, the Allen model includes a representation of the lateral geniculate nucleus (LGN), where arbitrary visual input is processed [12]. The resulting signals are relayed to V1 [12]. A more detailed description of the Allen model and its structure is provided in chapter 6.

1.3 Thesis Structure

Chapters 2 and 3 of this thesis present background material. An introduction into both the biology and the modelling methods relevant for this thesis is given. Chapter 2 describes the visual system of the mouse and the properties of nerve cells. In chapter 3, well-established options for modelling neurons are presented, including the two neuron representations which make up the Allen model: the Generalised Integrate-and-Fire neuron and the Linear-Nonlinear-Poisson neuron. Additionally, methods for modelling networks are explained. These are subsequently used to describe the structure and properties of the Allen network.

A theoretical foundation of neuronal responses is given in chapter 4 and 5. Chapter 4 introduces the spatial and temporal tuning properties explored in the Allen model. A mathematical approach is taken in chapter 5, which explains how neuron responses are described and modelled using linear filters. In this chapter, the linear response of a neuron to different spatial features is derived. The analytical results are later used to explain the responses measured in the LGN model.

The Allen model itself is described in chapter 6, partly by referring to the more general descriptions in chapters 3 and 5. In chapter 7, an overview of experimental findings of spatial and temporal tuning in the mouse LGN and V1 is presented. An evaluation of these simulation results is used to determine whether and how the model should be modified. The reader may use this chapter as an overview to look to for reference, and will find the most important findings listed in the provided tables and summaries there.

Chapter 8 deals with the method for creating visual inputs, running the Allen model, and extracting the tuning responses of the LGN and V1 cells. The method for making modifications to the model is also described. This chapter concerns primarily

technical details, while the motivation for the analysis will be given in chapter 9 and 10.

The findings from the Allen model simulations are presented in 9 and 10. Chapter 9 consists of multiple parts, where the first considers the original model. Subsequent sections present modifications made to the model and the resulting tuning curves. In all cases, findings are compared to experimental observations in mice. While the main focus of this thesis concerns the spatial properties in chapter 9, results of temporal tuning properties have also been included through chapter 10.

Finally, main findings and important key points will be summarised in chapter 11, followed by suggestions for further research.

Chapter 2

Neuroscience Background

2.1 The Mouse Visual Pathway

Out of the systems involving complex cortical processing, the visual pathway is maybe the most extensively studied [16]. Recently, a large portion of these studies have been performed on mice [16]. Using the mouse brain as a simplified model of the human brain has several advantages. Mouse physiology is, in many ways, similar to that of humans [17]. As much as 99% of the human genome is shared with mice [18]. In the visual cortex specifically, humans and mice share the same neuron subtypes in similar proportions, using a similar connectivity scheme [19]. It has also been argued that the relative simplicity of the mouse visual cortex makes it a more tractable model of the human brain, where general mechanisms can be more easily uncovered than in for example primates [19]. Furthermore, the mouse is small and practical to handle, easily available and inexpensive [20, 21]. Importantly, mouse cells can be easily manipulated using genetic tools [20]. Hence, the mouse is arguably an ideal subject for experimental researchers [18]. Therefore, much information now exists about the visual system of mice [20]. An illustration of the mouse visual pathway with its main components is provided in figure 2.1.

The eyes

As for humans, the first step in the visual pathway of the mouse is the photoreception [20, 22]. The process involves light being detected by rods and cones, which are photoreceptors on the retina [20]. These are cells which convert the light into electric nerve signals [20]. The signals are then filtered and shaped by retinal interneurons, which relay them further to the eyes' output neurons, called retinal ganglion cells [23]. In mice, like in other mammals, the retinal ganglion cells transmit to different structures in the brain, such as the lateral geniculate nucleus (LGN) in the thalamus and the superior colliculus in the midbrain [24, 25].

The lateral geniculate nucleus

The lateral geniculate nucleus (LGN) is a part of the visual system located in the thalamus [28] (see figure 2.1). It is the primary source of input to the visual cortex [25]. The LGN receives visual input directly from the retina, and relays the transformed nerve signals to the primary visual cortex, V1 [29]. The processing of the LGN involves spatial and temporal transformations, which modulate the input and decrease noise in the visual signal [30,31]. Like in humans, the mouse LGN is located between the retina

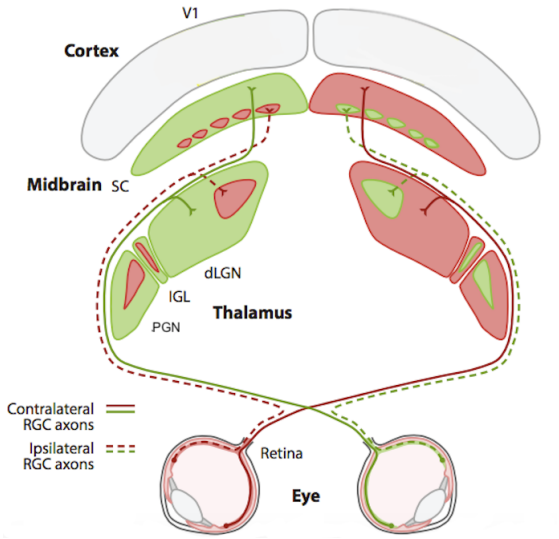

Figure 2.1: Illustration of the mouse visual system, from the retina in the eye to the primary visual cortex (V1). The three nuclei in the thalamus are the dorsal lateral geniculate nucleus (dLGN), the intergeniculate leaflet (IGL) and the pregeniculate nucleus (PGN). The superior colliculus (SC) lies in the midbrain. The two visual pathways each have one LGN and one SC [26, 27]. The contralateral and ipsilateral RGC axons refer to "opposite side" and "same side" axons from retinal ganglion cells. Adapted from Figure 1b in [20].

and the primary visual cortex [31]. The lateral geniculate nucleus (LGN) in mice is divided into the dorsal lateral geniculate nucleus (dLGN), the pregeniculate nucleus (PGN, previously called the ventral lateral geniculate nucleus) and the intergeniculate leaflet (IGL) [23, 24]. All three neuclei receive signals from the retina [23]. The dorsal lateral geniculate nucleus is the largest part of the LGN, and the main transmitter of visual signals for image formation in the brain [23]. It is also the area out of the three which is most extensively studied [23]. A lot less is known about the functions of the pregeniculate nucleus and intergeniculate leaflet [32]. Hence, dLGN is often simply referred to as LGN [33–35].

The primary visual cortex

The sensory cortex in mammals is organised into six layers (LI-LVI) which differ both functionally and anatomically [36, 37]. The topmost layer, LI, has few cell bodies, and consists primarily of axons and dendrites from deeper layers [37]. Layer II and III are referred to as the external granule and external pyramidal layer respectively, due to the regions being predominantly composed of granule and pyramidal cells [37]. In mice, the second and third layer cannot be distinguished in several areas, and are therefore referred to collectively as LII/III [38]. The internal granular cell layer, LIV, also contains granule cells. Typically, visual input first reaches the fourth layer [36]. This is the main target layer for feedforward connections [39]. The deepest layers are the internal pyramidal layer, LV, and the multiform layer, LVI [40]. The fifth layer contains pyramidal cells [40], while neurons in the sixth and most dorsal layer are of various types, such as pyramidal and fusiform cells [41] [37]. Both layers relay signals to subcortical areas [37].

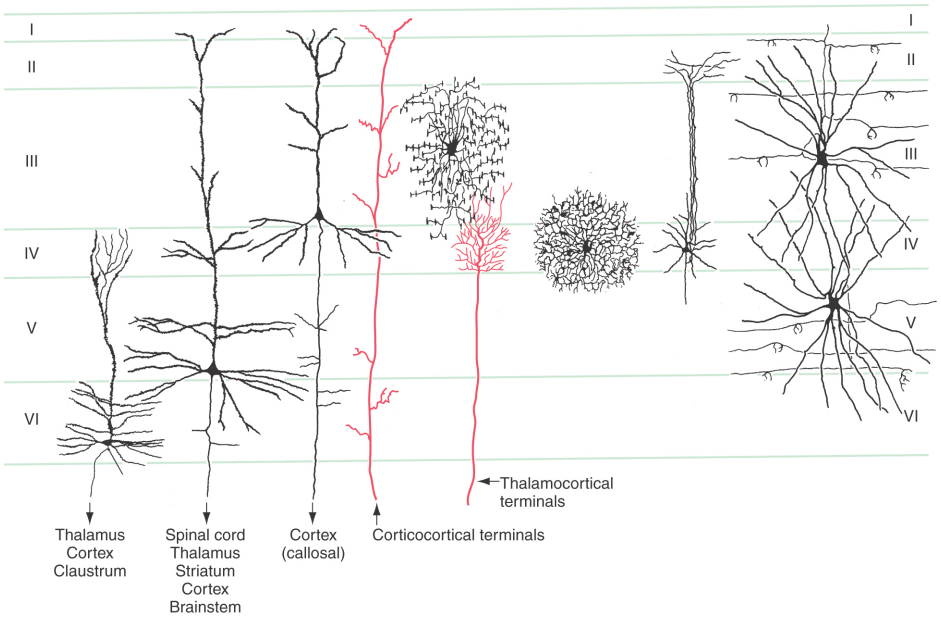

Figure 2.2: Example celltypes across layers in the cerebral cortex. The bottom labels note the destination or origin of the neural signals, illustrated by the arrows. Adapted from Figure 32-4. in [41, p.445].

Included in the sensory cortex is the visual cortex, which consists of different areas with distinct anatomy and properties [42]. Among these, the primary visual cortex (V1) is, in mice, the largest [39]. For most mammals, the neurons in V1 respond to complex features such as orientation of features, motion direction, stimulus depth (also called binocular disparity), as well as simpler features such as size [43]. The efferents of the mammal V1 are mainly to other areas in the visual cortex, but also back to the LGN or to other areas in the visual pathway [37]. An example is the feedback sent to the LGN from layer VI of the primary visual cortex [44]. Experimental evidence points to feedback inputs and interactions within the cortical network as necessary mechanisms for perception of complex images [45]. This can be ascribed to feedback signals carrying contextual information [46]. Neurons may therefore adapt their function depending on influences related to experience, expectation, attention or the task being performed [46]. One of the differences between the human and mouse V1 is the columnar organisation which is found in humans, but lacking in mice. Each column in the human V1 has a preference for specific visual features, such as orientation, color and direction [37]. Although the same structure is not found in mice, research shows that their neuron architecture is still organised to some degree, where neurons with similar orientation preference are weakly clustered together in smaller structures, so-called minicolumns [47]. Furthermore, similar to in humans, reproducible retinotopic maps of the mouse visual cortex has successfully been created [48].

2.2 The Neuron - Structure and Properties

Sections 2.2.1 and 2.2.2 are based on the books *Neuronal Dynamics* by Gerstner, Kistler, Naud and Paninski [9], *Computational Neurogenetic Modeling* by Benuskova and Kasabov [14] and *Principles of Neural Science* by Kandel, Schwartz and Jessell [25].

2.2.1 The nerve cell

The neuron consists of three main parts: the soma, axon and dendrites. A dendrite receives input from another neuron via the synapse. The neuron which sends the signal is called the presynaptic neuron, while the receiving neuron is called the postsynaptic neuron. When a nerve signal from a presynaptic neuron reaches the synapse, the presynaptic neuron sends out neurotransmitters: molecules which diffuse over to the postsynaptic neuron and cause the opening specific ion channels, which in turn enables current to flow through the neuron membrane. This causes a voltage response in the postsynaptic neuron, and the electric signal is transmitted towards its soma. Provided enough input reaches the soma, an action potential is initiated. The action potential travels down the axon, which is connected to the dendrites of subsequent neurons. In this manner, electric nerve impulses travel from neuron to neuron within the brain. As most action potentials, or spikes, have a similar shape, the information is conveyed by the timing and frequency of the spikes.

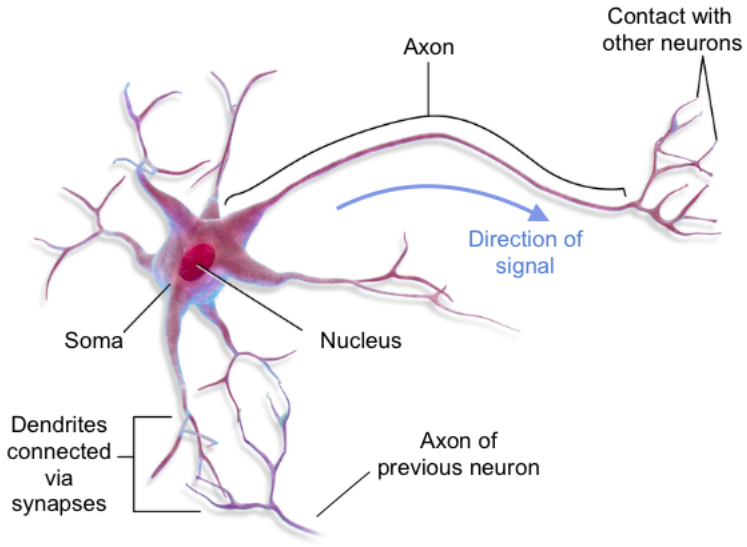

Figure 2.3: An illustration of the neuron structure. Adapted from [49].

The state of a neuron is usually described in terms of its membrane potential, that is, the voltage difference between the inside and outside of the cell. In the absence of input, the neuron has a specific, negatively polarised membrane potential ($\sim$ – 65mV [9]), called the resting potential. When a neuron receives synaptic input, the voltage changes. Synapses may be excitatory or inhibitory, which transmit positive or negative electric signals respectively. A neuron can receive multiple excitatory and inhibitory inputs from different synapses. Excitatory potentials cause depolarisation of the neuron, which if sufficient, causes an action potential to be generated. The inhibitory potentials have the opposite effect, as they prevent the neuron from spiking

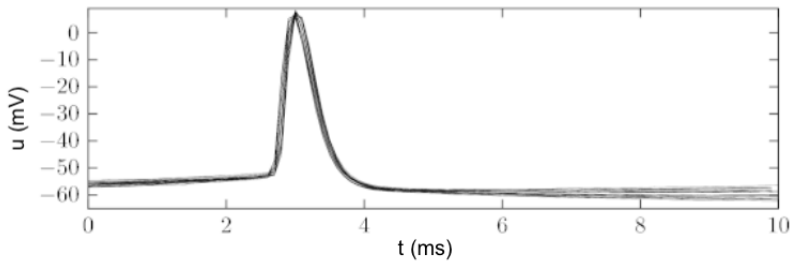

Figure 2.4: Recordings of several membrane potentials of a single neuron, aligned at the peak of the action potential. An action potential has a characteristic shape, with a rapid depolarisation, followed by a slower return to resting potential. Adapted from Figure 1.3 in [9].

by inducing further depolarisation. The inhibitory input is therefore referred to as being hyperpolarising. It is the total contribution from these inputs which decides whether the neuron fires or not.

2.2.2 Receptive fields and center-surround organisation

Collectively, the neurons in the brain are able to respond to a large repertoire of external stimuli. However, each neuron may only respond to a very specific stimulus. For example, in V1, neurons only respond to input from a confined region of visual space. This region is called the receptive field of the neuron. Recent descriptions sometimes also include a temporal component in this term [50], which will be introduced in section 4.3. For now, the more traditional definition [9, 25, 51] concerning spatial features is considered.

The ganglion cells in the eye have approximately circular receptive fields, which correspond to an area on the retina [25]. Most fields are composed of two features: the center and the surround, as illustrated by figure 2.5. The center and surround areas respond oppositely to light. The ganglion cells can be divided into two types based on their response. For on-center cells, bright stimuli excites the center area, while it inhibits the surround. In contrast, off-center cells have center regions that are inhibited by bright stimuli, while the surround is activated by light. If both regions are stimulated by the same light intensity, for example by a isoluminant screen, the two contributions are almost or totally cancelled out, and little response is generated by the cells. This structure is referred to as center-surround organisation, and not only found in ganglion cells, but in various cells along the visual pathway [25].

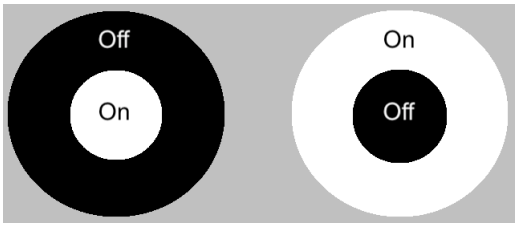

Figure 2.5: Illustration of circular center-surround receptive fields: on-center (left) and off-center (right). Based on the description by [25].

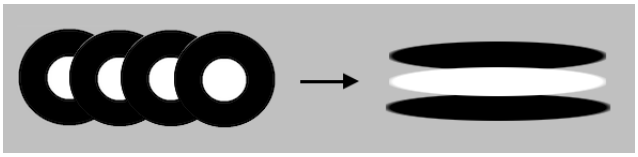

Figure 2.6: The receptive field of a simple V1 neuron (right) may arise from the receptive fields of the LGN neurons projecting to it (left). Based on the idea and description by Hubel and Wiesel [53].

The same center-surround structure of receptive fields have been found in the LGN of various animals, such as cats, monkeys and mice [25, 52]. This includes on-center and off-center neurons, and little response from cells to uniform, full field illumination [25]. The strong correspondence between the receptive fields of ganglion and LGN cells is partly due to only a small amount of ganglion cells projecting to each LGN neuron, so that the latter's receptive field structure is inherited from the ganglion cells [25]. In section 7, the experimental evidence for center-surround organisation in the LGN of mice will be considered in more detail.

In V1, the receptive fields are qualitatively different to previously in the visual pathway. While neurons in the retina and LGN typically prefer circularly symmetric stimuli, many neurons in V1 prefer elongated stimuli. The neurons of V1 can be classified as either simple or complex cells. The former, like LGN and ganglion cells, also have receptive fields with on and off regions. The latter often have elongated, elliptical receptive fields. The elliptical form of V1 receptive fields is believed to arise from the combination of LGN inputs transmitted to each V1 cell [25, 53]. Figure 2.6 illustrates this idea. For V1 cells with elliptical fields, the strongest response is produced when the stimuli is aligned along the axis of the receptive field. Perpendicularly oriented stimuli produces almost no response. The complex cells in V1 usually have larger receptive fields, and have no particular preference for stimulus position. Instead, stimulus orientation is the main determining factor for firing activity. Therefore, a favourable stimulus for both single and complex V1 cells is typically a moving bar, oriented along the cell's preferred axes. As a result, experimentalists often use moving bars or grating stimuli in order to produce the maximum neuron response in V1 [9].

Classically, the LGN has been thought to only comprise of simple centre-surround receptive fields, while the more complex mechanisms are found in V1. However, recent studies on mice show that properties such as preference for direction and orientation, as well as inhibition induced by visual contrast, is also found in the LGN [52]. Nevertheless, later findings suggest that direction selectivity is generated anew in the mouse visual cortex through integration of input with varying time delays and spatial locations [54].

Chapter 3

Computational Neuroscience

For more detailed descriptions of neuronal modelling, see the books *Neuronal Dynamics* by Gerstner, Kistler, Naud and Paninski [9] and *Principles of Computational Modelling in Neuroscience* by Sterratt, Graham, Gillies and Willshaw [13], which provide the basis for the following chapter.

3.1 Mechanistic Neuron Modelling

There are several different ways of modelling neurons: from simple models with a single dynamic variable to complex, biologically detailed models. In some situations, simple models may perform equally well as complex ones. Heavily detailed models may give more accurate results, but they are also less computationally efficient. Not to mention, complex models are often more difficult to tune, and having additional parameters is not necessarily an advantage if we do not know their function and how to optimise them. The choice of neuron model depends primarily on the type of stimuli, computational resources and modelling aim. For example, it is possible to model the exact shape and time development of the action potential. However, if the aim is to look at spike-times of a neuron, this may be an unnecessary level of detail.

In this section, different types of mechanistic neuron models will be introduced. These models take input in the form of synaptic current and describe the membrane potential dynamics which accounts for neuronal firing. First, the Hodgkin-Huxley model will be introduced, a description marking a breakthrough in numerical modelling of neuronal dynamics. It will then be explained how this relatively complex model can be reduced into a Leaky Integrate-and-Fire neuron. Subsequently, the Generalised Leaky Integrate-and-Fire model (GLIF) will be described. Various GLIF models are the building blocks of the V1 Allen model [12], the neuronal network which are explored in this thesis. Lastly, for the sake of completeness, options for modelling the spatial structure of neurons will be mentioned.

3.1.1 The Hodgkin-Huxley model

Hodgkin and Huxley are often viewed as the first to successfully predict results in experimental neuroscience using numerical methods. Their model describes the ion fluxes and voltage changes in the neuron and how this gives rise to, and shapes the action potential. The original Hodgkin-Huxley model includes three key types of ion channels in the neuron membrane: the sodium, potassium and "leak" channel. The "leak" channel correspond to the remaining, mainly chlorine, ions. The current

flowing through these three channels govern the cell's voltage dynamics and generate the action potential. The dynamics can be expressed as an electric circuit [9]

$$C\frac{du}{dt} = I(t) - \sum_i I_i(t) \tag{3.1}$$

where C is the capacitance of the cell membrane, u is the membrane potential, $I(t)$ is the synaptic input current to the neuron and $I_i(t)$ is the current through the ion channels of type i. The current through the ion channels is governed by [9]

$$\sum_i I_i = g_{Na}m^3h(u - E_{Na}) + g_k n^4(u - E_k) + g_L(u - E_L) \tag{3.2}$$

where g_{Na}, g_k and g_L are the conductances and E_{Na}, E_k and E_L the reversal potentials of the sodium, potassium and leak channels respectively. The reversal potential is the membrane potential at which the flow of a specific ion across the cell membrane changes direction [13]. The "gating variables" m, h and n determine the probability of a channel being open. They depend on the membrane potential via [9]

$$\frac{dx}{dt} = -\frac{1}{\tau_x(u)}(x - x_0(u)) \tag{3.3}$$

where x represents m, n or h. In other words, each gating variable approaches its own steady state value, $x_0(u)$, with a specific time constant $\tau_x(u)$ [9]. The Hodgkin-Huxley model can be extended to incorporate other types of ion channels, so that it may describe various types of neurons, synapses and more complex dynamics.

Following the Hodgkin-Huxley model, a large number of neuron models with different levels of complexity have been established. One limitation of the Hodgkin-Huxley model is that it is relatively complex. However, by approximating the evolution of the gating variables by a single equation, the model is reduced to two dynamic variables: the membrane potential and one effective gating variable. The model can then be both visualised and analysed analytically. This decreases the computational load and makes it easier to analyse certain aspects of neuron behaviour.

Going a step further, one can exploit the difference in time scales at which the voltage and effective gating variable evolve. In the subthreshold regime, the gating variables evolve slowly compared to the membrane potential. As long as the potential is below threshold, one can therefore approximate the gating variables by a constant value. As soon as the potential reaches the threshold, the spike itself can be replaced by a simple reset of the voltage. We then arrive at the classic Integrate-and-Fire neuron.

3.1.2 The Generalised Leaky Integrate-and-Fire model

In the Integrate-and-Fire model, one takes advantage of the fact that the action potential has a characteristic voltage time development. When neurons generate action potentials, the main information is contained in the timing of these spikes. The Integrate-and-Fire model therefore reduces action potentials to events in time. The simplest form of this model is called the Leaky Integrate-and-Fire model (LIF). In this case, the spike events are modelled as a reset of the voltage to resting potential whenever the membrane potential reaches a threshold. That is [9],

$$\text{if } u(t) = u_{\text{threshold}} \quad \text{then } \lim_{\delta t \to 0, \delta t > 0} u(t + \delta t) = u_{\text{rest}}. \tag{3.4}$$

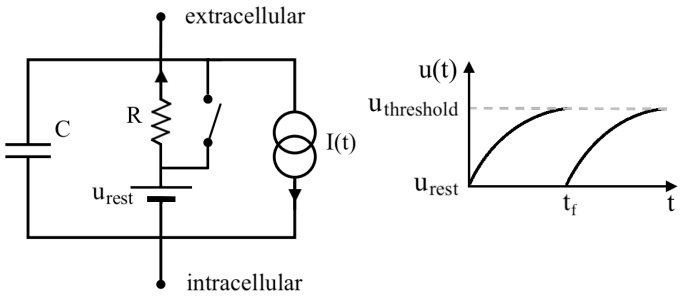

(a) (b)

Figure 3.1: (a) The dynamics of an Integrate-and-Fire neuron, as defined by equations (3.4) and (3.5), can be represented by an electrical circuit. The switch closes when the membrane potential reaches the firing threshold, so that the potential is reset to u_{rest}. Inspired by Figure 8.9a in [13, p.199]. (b) The membrane potential as a function of time in response to a constant example input current. Given sufficiently strong input, the membrane potential reaches the threshold at time t_f, at which point the neuron fires and is reset to u_{rest}.

where u is the membrane potential of the neuron. The voltage between the firing times is governed by a single equation [9],

$$\tau \frac{d\,u}{dt} = -(u(t) - u_{\text{rest}}) + RI(t), \tag{3.5}$$

where τ is the membrane time constant, R is the membrane resistance and $I(t)$ is the synaptic input current. The equation above describes the charge leaking through the neuron membrane. In order to make the model more realistic, one can for example add a dynamic threshold or after-spike currents.

In the case of adding after-spike currents, spike induced effects over relatively long timescales are taken into account. As described in section 3.1.1, the initiation of a spike causes certain ion channels to open. In addition to forming the rapid and stereotypical action potential, some ion channels may have slower effects on the membrane potential. These effects can be modelled as currents, activated by the spiking of the neuron. An extra term is then added to equation (3.5), so that it becomes [1]

$$\tau \frac{d\,u}{dt} = -(u(t) - u_{\text{rest}}) + RI(t) + R\sum_j I_{\text{j}}(t), \tag{3.6}$$

where $\sum_j I_{\text{j}}(t)$ is the sum of the after-spike currents. Each time the neuron spikes, the value of the current is incremented by a_j. Specifically [1],

$$\text{if } u(t) = u_{\text{threshold}} \quad \text{then } \lim_{\delta t \to 0, \delta t > 0} I_{\text{j}}(t + \delta t) = I_{\text{j}}(t - \delta t) + a_j. \tag{3.7}$$

Between firing times, the after-spike current dynamics are governed by [1]

$$\frac{dI_{\text{j}}(t)}{dt} = -b_j I_{\text{j}}(t); \quad j = 1, ..., N. \tag{3.8}$$

The above description, defined by equations (3.6) - (3.8), has been used to model neurons in the Allen network [12, 55] which will be explored later (see chapter 6.1).

By adding after-spike currents or other modifications, while still maintaining a linear voltage equation, the LIF model is extended to a Generalised Leaky Integrate-and-Fire (GLIF) neuron.

The GLIF neuron is an example of a model which, despite being very simple, may produce realistic results. For example, it accurately predicts the membrane potential and spikes in a real neuron when driven by a varying input current [9]. Even though the model is a substantial simplification of the real neuron, it agrees well with experimental observations and more biophysically detailed models [9–12]. In addition to providing a computationally efficient candidate for network simulations, the simplicity of the LIF and GLIF cell make the network as a whole easier to analyse [13].

3.1.3 Spatially extended models

While the above descriptions have focused on the generation of spikes in the soma, a neuron model may also include spatially extended structures. Dendrites can be represented as cylindrical segments which allow current to pass from synapses to the soma. Complicated morphologies can be constructed by having cylindrical segments of different diameter and length. Axons can be modelled in a similar way as dendrites. While dendrites usually have passive drift of current, the active properties of the axon can be reproduced using the Hodgkin-Huxley description. Models with cylindrical segments which represent the spatial structure of the neuron are called compartmental models. They allow for a high level of variety and biophysical detail. This enables us to reproduce distinct features of different neuron types. For example, biophysically detailed models based off experimental classifications in the Allen Cell Types Database[1] were able to exhibit properties characteristic of different cell types [56].

3.2 Statistical Neuron Modelling

The Hodgkin-Huxley and GLIF model in the form they are presented above do not take visual stimuli as input, but rather synaptic current. A fundamental problem in neuroscience is how stimuli is transformed from visual input to spike trains [57]. Models which have been used to describe this process are mainly statistical instead of mechanistic [58]. That is, instead of modelling the membrane potential dynamics of the neuron, the model describes the statistical relationship between stimuli and neuron activity [59]. This relationship is found by optimising a set of variables against experimental measurements [59]. Among the simplest statistical models for encoding stimuli and estimating the resultant spike train is the Linear-Nonlinear Poisson (LNP) model [60]. The LNP model is used to represent the LGN cells of the Allen network [12] which is explored later.

3.2.1 The Linear-Nonlinear-Poisson model

The Linear-Nonlinear Poisson (LNP) cascade was introduced as a relatively simple way of capturing neuron responses [60]. It is a different type of model to the Hodgkin-Huxley and Integrate-and-Fire neuron in the sense that it does not consider the synaptic input or voltage dynamics of the neuron. Instead, the LNP model gives a

[1]The Allen Cell Types Database is available at: `https://celltypes.brain-map.org/`

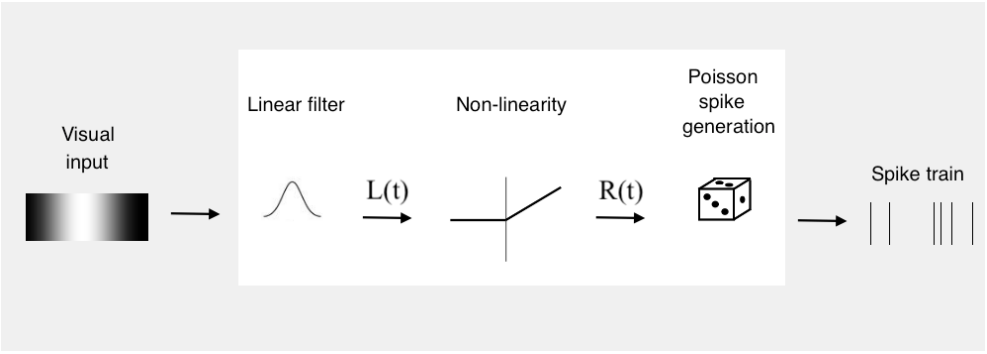

Figure 3.2: The LNP model takes visual stimuli as input. A linear filter is applied to the input, and the result is passed through a non-linear function. This gives the firing rate, $R(t)$. Spikes are then generated from the firing rate via a Poisson process.

statistical description of the neuron's response by taking the visual stimuli as input directly. In contrast to the previously mentioned models, the LNP model does not predict exact spike times, but instead produces stochastic spike times based on calculated firing rates.

The steps of the LNP cascade are as follows. The model first performs a linear operation on the stimulus. This operator may include both a spatial and a temporal component, which describes how different spatial and temporal stimulus features are processed [60]. The linear response of a neuron will be described mathematically in section 5. To avoid negative firing rates, it is necessary to introduce a non-linear transformation, for example in the form of a threshold [50]. To make the description general, the neuron activity in terms of its firing rate can be described by [60]

$$R(t) = N(L(t)) \tag{3.9}$$

where the function N is an arbitrary non-linear function and $L(t)$ is the linear response to the stimulus. One option for the non-linearity is to add a constant, spontaneous firing rate to the linear response, and pass both terms through a rectifying unit, so that the response becomes [12]

$$R(t) = \Gamma ReLU(f_{sp} + L(t)) \tag{3.10}$$

where Γ is a constant and f_{sp} is a spontaneous firing rate independent of the input. The Rectified Linear Unit function, $ReLU(x)$ [61], is given by

$$f(x) = \begin{cases} x, & \text{if } x \geq 0. \\ 0, & \text{otherwise.} \end{cases} \tag{3.11}$$

The non-linear function in equation (3.10) is the description used for the Linear-Nonlinear Poisson cells in the Allen LGN [12].

After calculating the firing rate of the neuron, the rate is converted to a spike train through a Poisson process [60]. This process generates spikes randomly with a probability defined by the instantaneous firing rate [60]. The model only uses visual stimuli to predict firing rates. It cannot account for influences such as recurrent connections between neurons [13] or previous spiking history [60]. The LNP model is therefore a highly simplified description of real neuron responses.

3.3 Modelling Networks

A neuronal network consists of neurons relaying and receiving spike input from each other through their synapses. Multiple issues have to be considered when modelling such a network. Some of these include choosing the appropriate neuron model, number of cells, type of connections and synapse description.

Simplifications and variability

When creating models of neuronal networks, several simplifications have to be made. While it can be reasonable to use detailed neuron models when looking at the behaviour of a single or a few cells, it is generally not feasible if the aim is to model a large network. The larger the network, the simpler the neuron model required for the simulations to be computationally efficient. Furthermore, if one wants to model a part of the cortex containing thousands of nerve cells, the number of cells may have to be scaled down and the connectivity between them appropriately adjusted.

Another aspect which needs to be addressed is the variability of cells. A network of cells may consist of identical models with the same parameters. Realistically, however, there will be some variation. This could be introduced into the model by for example sampling the neuron parameters from a distribution, or by basing them on experimental fits to different neuron types. Using a network of GLIF neurons as an example, the values of threshold, membrane potential, membrane time constant and resistance may vary between cells.

Connectivity

A major defining property of a neuronal network is the connections between cells. Whether the network is scaled down or full size, there are many options to choose from. While it is possible to have a so-called full connectivity, where every cell is connected to every other cell, it is more realistic to have each cell connected to only a fraction of the cells in the network. A common scheme is to randomly establish connections between cells with a probability depending on the inter-cell distance. One can for example use a Gaussian probability distribution [12],

$$P(r) = Ae^{-r^2/\sigma_r^2} \tag{3.12}$$

where r is the distance between cells, σ_r is the standard deviation and A is a scaling factor. The scaling factor influences the total number of connections established from a cell [12]. The probability could also depend on which cortical column or layer a neuron belongs to. Another option is to have the connectivity depend on the cells' stimuli preferences and whether they are inhibitory or excitatory. The connectivity between real neurons is not fully known, but there is experimental support for all the above suggestions [9, 12].

Synapse description

The synapses connecting the cells can be modelled in different ways. One option is to use a circuit description, where the current depends on the membrane voltage and conductivity as [9]

$$I_{syn}(t) = g_{\text{syn}}(t)(u(t) - E_{\text{syn}}) \tag{3.13}$$

where $I_{\text{syn}}(t)$ is the synaptic current, $g_{\text{syn}}(t)$ is the conductivity and E_{syn} is the reversal potential of the synaptic ion channels. An increase in conductivity represents

the effect of the neurotransmitters on the postsynaptic ion channels: As more ion channels open, more current is able to flow through the neuron membrane, which in turn generates a postsynaptic potential. The conductivity may depend on the incoming spike times in the form of an alpha function [13]

$$g_{\text{syn}}(t) = \bar{g}_{\text{syn}} \frac{t - t_f}{\tau_{\text{syn}}} e^{-(t-t_f)/\tau_{\text{syn}}} \, \Theta(t - t_f) \tag{3.14}$$

where $\bar{g}_{syn}$ is the maximum conductance, t_f is the arrival time of a spike, τ_{syn} is the synaptic time constant and $\Theta(t - t_f)$ is the Heaviside step function. The time dependence can also take similar forms such as a single exponential decay or a difference of exponentials. A simpler alternative to the conductance based synapse in equation (3.13) is to have the input current independent of membrane potential. This is called a current based synapse and can for example be expressed as [13]

$$I_{\text{syn}}(t) = \bar{I}_{\text{syn}} \frac{t - t_f}{\tau_{\text{syn}}} e^{-(t-t_f)/\tau_{\text{syn}}} \, \Theta(t - t_f). \tag{3.15}$$

where $\bar{I}_{\text{syn}}$ is the maximum input current. Implicit in $\bar{g}_{\text{syn}}$ or $\bar{I}_{\text{syn}}$ is the overall strength of the synapse, also called the synaptic weight. The weight is often a constant which is either the same for all connections or set based on factors such as distance and similarity between cells. It is also possible to use dynamic weights. One option is to include a dependence on spiking arrival history, as this may influence for example the availability of neurotransmitters.

Chapter 4

Spatial and Temporal Neuronal Response Properties

This thesis explores selected properties of the Allen model: size tuning, spatial frequency tuning and temporal frequency tuning. This chapter describes the different properties and how they can be measured using various visual stimuli. As will be seen in chapter 7, these are properties which are often investigated in real neurons; they help classify neurons and describe their responses to basic visual stimuli. In the following, cell qualities giving rise to different tuning behaviour will be presented. In this way, the chapter explains what observed responses tells us about cells' receptive fields (see section 2.2.2 on receptive fields), and hence provides a basis for how neurons can be represented mathematically. This mathematical representation will be the topic of chapter 5.

4.1 Size Tuning

Size tuning refers to a cell's response to stimuli as a function of stimuli size [62]. A central property related to size tuning is called "surround suppression" [63]. The latter refers to responses to stimuli centered in the receptive field being inhibited by stimuli in the surround region [63]. Surround suppression is thought to be partly responsible for the brain's ability to distinguish borders, object contours and precise line segments [25], by increasing the perceived contrast between an object and its background [64]. For cells with surround suppression, a stimuli which is centered in

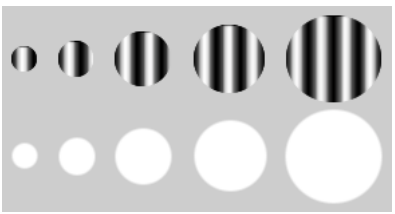

Figure 4.1: Example sets of stimuli for measuring the size tuning of a neuron. Patch gratings (top) and white spots (bottom) of increasing sizes.

the middle of a neurons receptive field can be used to measure suppression [63]. The activity will increase as long as the stimuli is contained within the "center" region [63]. After reaching a certain, preferred size, the firing rate will start decreasing as the stimuli extends into the "surround" area [63]. In this thesis, two main types of size tuning stimuli are used: flashing spots of uniform luminosity and patch gratings, both illustrated in figure 4.1. The second type of stimuli, patch gratings, consist of a circular patch of a horizontally moving sinusoidal grating.

4.1.1 Linear and nonlinear surround suppression

Suppressive mechanisms in neuronal networks are usually divided into two categories: linear (classical) and nonlinear (extraclassical) surround suppression [65]. While linear suppression is relatively simple and can be explained by the structure of a neuron's receptive field, nonlinear suppression is a collective term used to describe suppressive effects this does not account for [63, 66, 67].

Linear suppression

Section 2.2.2 explains the on-off antagonism of receptive fields, in the form of circular center-surround cells in the LGN and either circular center-surround or adjacent elliptical subfields in V1. These structures may give rise to what is referred to as classical or linear surround suppression [65]. From linear algebra we know that a transformation is linear if the following conditions apply [68]

$$f(x_1 + x_2) = f(x_1) + f(x_2) \tag{4.1}$$

$$f(\lambda x) = \lambda f(x). \tag{4.2}$$

As will be seen in chapter 5.2, the center-surround structure of a receptive field can be described by a linear function of the stimulus luminosity. In this context, x_1 and x_2 represent the stimuli, f is the receptive field response and λ is a scaling factor. Whether linear surround suppression occurs, depends on the type of neuron and stimuli presented [63]. In general, a linear on-center cell will display surround suppression in response to bright stimuli, and an off-center cell will show surround suppression for dark stimuli [25]. In both cases, the neuron activity increases if said stimuli covers more and more of the receptive field center, before decreasing as the stimuli extends into the surround area [25].

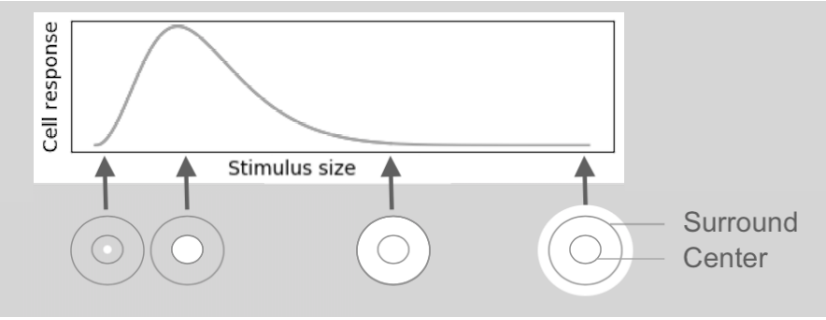

Figure 4.2: The response of an on-center neuron to bright circular stimuli depends on stimuli size. The white circles represent the visual input covering various portions of a cell's center-surround receptive field.

An example of the response of an on-center cell to bright stimuli is given in figure 2.5. The response to more complex stimuli will be discussed in section 5.2. Note that in reality, it is not possible for a cell to behave completely linearly [9]. At minimum, there has to be some form of rectification of the linear response in order to avoid negative firing rates [9]. Nevertheless, suppression which can be explained by this kind of cell is still referred to as linear suppression, since it arises from the on-off antagonism of the linear receptive field [63] [65].

Nonlinear suppression

Neuronal responses which cannot be explained by a linear receptive field, plus a nonlinear saturation or threshold, have been observed in the visual pathway of various animals [65,67,69]. This involves modulatory effects from regions outside the classical receptive field, as well as nonlinear responses within the classical receptive field (apart from a simple threshold or response saturation) [67]. These properties have been given the collective name of extraclassical (or nonlinear) effects [67]. One such effect is the surround inhibition which is not accounted for by the antagonistic center-surround structure [63, 66, 70].

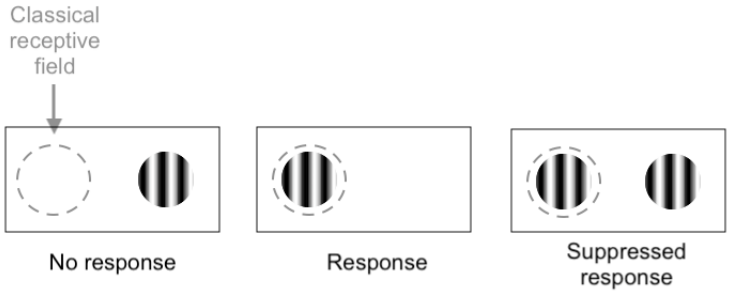

Figure 4.3: An example of nonlinear surround suppression. Stimuli in the extraclassical surround does not by itself generate a response [71] (left). It may, however, suppress the response to a stimulus which is inside the classical receptive field [71] (right).

The classical, center-surround receptive field in section 2.2.2 refers to the region where stimuli can produce neuronal firing. Stimuli outside this region alone does not evoke neuronal activity, but can have a modulatory effect on stimuli within the classical receptive field [72]. The outside region is a part of the extraclassical receptive field [71]. The existence of extraclassical receptive fields have been reported in the retina, LGN and visual cortex of different animals [65] [69]. Stimuli in the extraclassical surround can have either a facilitatory or inhibitory effect [72], but have usually been found to be suppressive [65, 70, 71, 73]. For example, in mice V1, the response to a patch grating within the classical receptive field is suppressed by a grating in the extraclassical surround [64]. Furthermore, this suppression is found to be orientation dependent [64]. That is, the suppression is stronger if the surround grating has the same orientation as the center grating than if it is cross-oriented [64] (see figure 4.4).

For stimuli contained within the classical receptive field, a number of properties separates extraclassical and classical suppression. In contrast to the on- and off-regions of the center-surround receptive fields, nonlinear suppression may inhibit activity for both bright and dark stimuli [66]. This gives rise to size inhibition in response to stimuli of varying luminosity, such as patch gratings, which is not fully accounted for by a linear receptive field [63, 71]. For example, by fitting the receptive

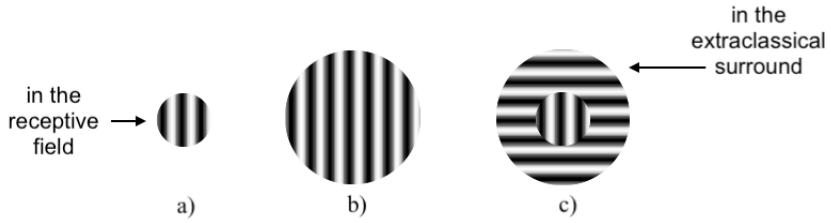

Figure 4.4: The response to a patch grating presented within the classical receptive field can be suppressed by stimuli in the extraclassical receptive field [64]. The center patch grating is the same in all three cases [64]. The suppression is stronger if the surround grating is oriented along the same direction as the center (b), as opposed to being cross-oriented (c) [64]. The illustration is inspired by results in [64].

fields of cat LGN cells to a linear model, Alitto and Usrey [63] found that the predicted suppression caused by the center-surround structure was substantially weaker than what was observed. Another effect which indicates the presence of a nonlinear surround is that the strength of suppression depends on stimuli contrast [63, 70].

The origin of extraclassical surround suppression in the visual pathway is not yet properly understood [65]. Many explanations have been proposed to account for the phenomena, and past studies indicate that various mechanisms may be involved. For example, one study suggests that neuronal circuits in various V1 layers may give rise to extraclassical surround suppression in V1, as the orientation and contrast tuning of the suppression varied across cortical layers [74]. Another experiment argues that surround suppression in V1 may primarily be due to reduced input from the LGN [62]. Different delays in onset of suppression suggests the presence of at least two different mechanisms: One may originate from feedforward inputs to V1 and another from recurrent V1 connections or feedback from other cortical areas [75].

For extraclassical suppression in LGN, some findings suggest that cortical feedback is involved, as cortical inactivation reduced the suppression in LGN [76–79]. An experiment on mice LGN and V1 also showed that stimulation of cells in V1 suppressed neuronal activity in LGN [80]. Others point to inhibitory recurrent connections in LGN as possible sources [81]. A study on the macaque monkey concluded that nonlinear suppression in the LGN relies on feedforward projections from the retina [66]. Their argument partly rests on there being no significant difference of suppression strength in the retina versus LGN [66]. Additionally, they report that suppression in LGN arose too quickly to be the result of cortical feedback [66]. A later study showed that about 70 % of the surround suppression in the cat LGN was also displayed in retinal cells [63]. Their results imply that feedforward signals play a major role in extraclassical surround suppression, and that it first originates in the retina before being amplified and modified in the LGN and V1 [63]. In conclusion, several mechanisms appear to be involved in surround suppression of stimuli.

4.1.2 Suppression index

When analysing the dependence of neural responses on size in the Allen model, the degree of suppression will be assessed qualitatively, by constructing the tuning curves of each cell. All tuning curves will also be included for reference. However, when looking at previous experiments, a suppression index (SI) is sometimes used to quantify and present the degree of recorded suppression [64, 82, 83]. In chapter 7, the suppression index will therefore be used to refer to certain experimental findings. Where defined

explicitly, the suppression index used in these studies has been given as [64, 82]

$$\mathrm{SI} = \frac{R_{\mathrm{peak}} - R_{\mathrm{large}}}{R_{\mathrm{peak}}} \tag{4.3}$$

where R_{peak} is the activity at peak response and R_{large} is the activity at the largest stimuli. The value of SI ranges from 0 to 1. A suppression index of 0 corresponds

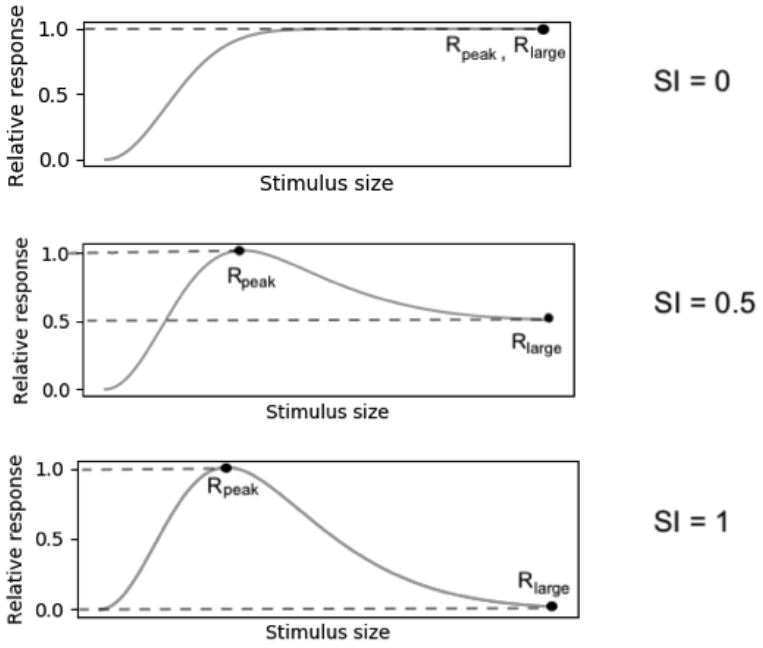

Figure 4.5: Size tuning for three hypothetical cells with different surround suppression index. From no suppression (SI = 0) to maximum suppression (SI = 1).

to no surround suppression, where the largest stimuli gives peak response. For SI equal to 1 there is maximal suppression, with no response to the largest stimuli. The response of three example cells with SI = 0, 0.5 and 1 are shown in figure 4.5.

4.2 Spatial Frequency Tuning

Recording a neuron's spatial frequency tuning has historically been a popular way of assessing its response properties [84]. Besides being relatively straightforward to measure, spatial frequency tuning can tell us much about the spatial structure of a cell's receptive field [85]. Studying a neurons receptive field via its spatial frequency tuning is called "spatial frequency analysis" [85]. It involves presenting a sinusoidal grating of different spatial frequencies as visual stimuli to the neuron [85]. Experiments have shown that for a cell which behaves like a linear filter, grating stimuli can be used to map its receptive field [85]. Furthermore, responses to any stimuli, such as flashing spots, can be predicted from the response of a linear cell to grating stimuli [85–87]. It is therefore possible to predict the size tuning of the cell and whether it will exhibit linear surround suppression. The relationship between the frequency tuning curves and surround suppression will be explained mathematically in chapter 5.2.

Figure 4.6 illustrates how the optimal spatial frequency (and orientation) of a grating stimulus depends on the structure of the receptive field. Center-surround cells (figure 4.6 b) are typical in the retina and LGN, while elliptical on- and off-regions (figure 4.6 c,d) are often seen in V1 [25].

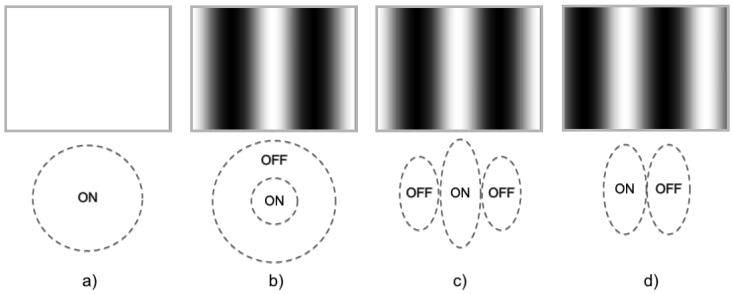

Figure 4.6: Gratings with optimal spatial frequency (top) for a selection of receptive fields (bottom). (a) A purely excitatory field gives largest response for uniform bright stimuli (zero spatial frequency) [88]. (b) An on-center receptive field has a non-zero preferred spatial frequency [87]. (c, d) When the bars of the grating is aligned along the orientation of the receptive fields, the elliptical on- and off-regions give rise to a non-zero preferred spatial frequency [89].

A cell which is purely excitatory or inhibitory prefers stimuli which are uniform or gratings of very low spatial frequencies, resulting in a so-called low-pass frequency tuning curve [88] (see figure 4.7 b). By contrast, cells with center-surround organisation have spatial frequency tuning curves with a band-pass shape, with a peak response to a non-zero frequency [87] (see example in figure 4.7 a). Receptive fields which consists of elliptical on- and off-regions can also give rise to band-pass tuning [89]. The larger the number of alternating on- and off-regions, the narrower the bandwidth of the frequency tuning [89]. Whether the receptive field of a cell consists of a center-surround circle or alternating elliptical subfields, the band-pass tuning arises from the antagonism between the excitatory and inhibitory regions [89, 90].

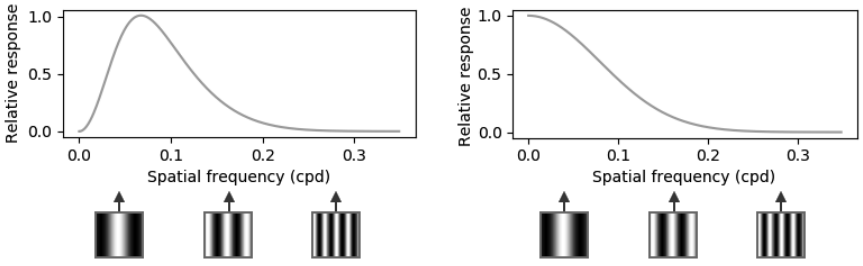

(a) Band-pass filter

(b) Low-pass filter

Figure 4.7: Two types of spatial frequency tuning. The illustrations below the graphs show sinusoidal gratings of three different spatial frequencies. Spatial frequencies are given in cycles per degree (cpd).

Spatial frequency analysis has been used extensively as a method for analysing cells' visual receptive fields for the last 50 years [85]. The method has been applied

to cells found in areas such as the retina, LGN and visual cortex [85]. Although the visual pathway also contain nonlinear cells, the method provides a good starting point for understanding cells in the visual pathway [91]. Practical reasons make spatial frequency tuning curves a good alternative to size-tuning curves when mapping neurons receptive fields. In particular, the full field grating stimuli does not require knowledge about the visual center of a neuron [85]. In addition, no fine tuning of size or use of very small stimuli is required, which makes it feasible to investigate even very small receptive fields [85].

4.3 Temporal Frequency Tuning

The visual receptive field of a neuron was originally used to describe the region in visual space for which the neuron is sensitive [51]. This is likely still the most common definition [9], and concerns only the spatial tuning properties of the neuron. However, recently the term has been expanded to include temporal selectivity, in a so-called spatio-temporal receptive field [50]. Following this definition, one can think of temporal frequency analysis as the analogue to the spatial frequency analysis presented above.

As for spatial tuning, the preferred temporal frequency of a neuron depends on the structure of its temporal receptive field [50]. Neurons typically prefer stimuli which changes with time, such as drifting gratings and flashing spots, as opposed to static gratings or spots [50]. This preference for non-zero temporal frequencies again arises from antagonisms in the receptive fields, but now in the temporal component [50]. To illustrate this point, figure 4.8 shows a typical temporal response of a receptive field, where the optimal luminosity has been superimposed on the receptive field function. The illustration shows the response of the field as a function of time since stimulus,

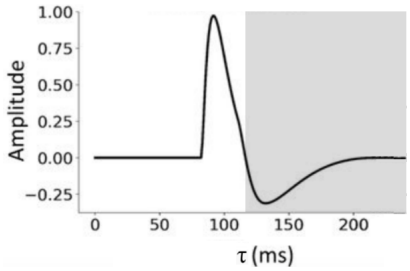

Figure 4.8: The temporal receptive field of a neuron (represented by the solid line) is typically bi-phasic [50], giving rise to a preference for stimuli which changes with time (background shading). The temporal field function is equal to the response of the field as a function of time, τ, since showing a delta pulse of light [50]. Adapted from Figure 2 in [12].

τ. The response is so-called bi-phasic [92], were the field has a change in polarity over time. Hence, the maximal response will be evoked when a dark stimuli switches to a bright one [50]. Note that the value of τ increases as we consider stimuli further back in time. This preference is assuming the contribution from the spatial field is positive, such as the on-cell in figure 4.6a. If the same temporal structure is found in for example an off-cell, the shading in figure 4.8 would be opposite.

In summary, temporal receptive fields with changes in polarity will, analogous to the spatial fields, give rise to band-pass temporal frequency tuning, as opposed to the low-pass tuning which would be produced by a simply positive field.

Chapter 5

Modelling Visual Receptive Fields

5.1 Separable Receptive Fields

Although a cell model must include a non-linear rectification to prevent negative firing rates, a linear description goes a long way in explaining neuronal properties [50]. A neuron's response to spatial and temporal features of visual stimuli may be referred to as its spatio-temporal receptive field. The linear response of such a receptive field can be written as an integral [50]

$$L(t) = \int_0^\infty \int\int d\tau dx dy D(x, y, \tau) S(x, y, t - \tau) \tag{5.1}$$

where $D(x, y, \tau)$ describes the effect of a stimulus at the point (x, y) and time $t - \tau$ on the firing rate at time t. The parameter τ hence refers to time since stimulus. $L(t)$ gives the neuron's response in terms of its firing rate. If the filter is spatio-temporally separable, $D(x, y, \tau)$ can be expressed as a product of the spatial and temporal part, so that the linear response is given by [50]

$$L(t) = \int_0^\infty \int\int d\tau dx dy D_s(x, y) D_t(\tau) S(x, y, t - \tau) \tag{5.2}$$

where $D_s(x, y)$ is the spatial kernel, $D_t(\tau)$ is the temporal kernel and $S(x, y, t - \tau)$ is the stimulus. Furthermore, given that the stimulus can be written as $S(x, y, t - \tau) = S_s(x, y)S_t(t - \tau)$, the integral becomes [50]

$$\begin{aligned} L(t) &= \int_0^\infty d\tau D_t(\tau) S_t(t - \tau) \int\int dx dy D_s(x, y) S_s(x, y) \\ &= L_t(t) L_s, \end{aligned} \tag{5.3}$$

where L_s and $L_t(t)$ describe the spatial and temporal response components respectively.

The spatial kernel corresponds to the spatial receptive field. A simple example is the circular 2D Gaussian function [12],

$$D_s(x, y) = \frac{1}{2\pi\sigma^2} e^{-(x^2+y^2)/2\sigma^2} \tag{5.4}$$

were σ is the standard deviation. A positive Gaussian represents an on-region, which gives a positive response for bright stimuli [12]. Conversely, a negative Gaussian describes an off-region [12]. A variety of spatial kernels have been used to model receptive fields in neurons, such as the Difference-of-Gaussians [93] and the Gabor filter, [94] which will be described in sections 5.2 and 5.4 respectively.

The temporal kernel defines the effect of previous visual input, at time $t - \tau$, on the activity at time t [50]. This may for example be expressed by a bi-phasic function [12, 92]

$$D_t(\tau) = w_1 \frac{\cos(\log(t_1 + \tau) - d_1) + 1}{2} + w_2 \frac{\cos(\log(t_2 + \tau) - d_2) + 1}{2} \tag{5.5}$$

where t_1 and t_2 are time constants, w_1 and w_2 are weights and d_1 and d_2 are offsets. An illustration of this function, where w_1 and w_2 have opposite polarity, has been shown in figure 4.8 in section 4.3.

The two kernels given by equation (5.4) and 5.5 constitute the spatio-temporally separable filter used by the Allen LGN model [12]. Since the complete kernel is separable, the response of the spatial filter can be analysed separately from the temporal filter response [50].

5.2 The Difference-of-Gaussians Filter

The Difference-of-Gaussians (DoG) model was established by Rodieck in 1965 [93]. The function has subsequently been used to describe spatial receptive fields and responses of cells in the LGN, retina and, although less so, in the visual cortex of different species [76, 95]. The DoG model describes receptive fields with a center-surround structure and linear surround inhibition [96]. The filter is composed of two overlapping, concentric 2D Gaussian functions of opposite polarity [93]. An on-center/off-surround field is represented by a narrower, positive center Gaussian and a wider, negative surround Gaussian with a smaller amplitude [93]. An illustration of

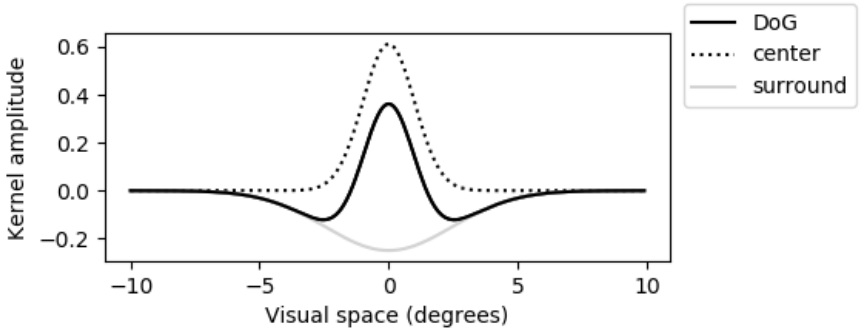

Figure 5.1: Illustration of the on-center/off-surround DoG filter given by equation (5.6). The center and surround terms are shown by the dotted and gray line respectively, while the sum of the two, the resultant DoG filter, is shown by the solid black line. Parameters are $A_c = A_s = 1.73$, $\sigma_c = 2$ and $\sigma_s = 2.45\sigma_c$. Values for σ and σ_c refer to a representative cell in the LGN model, while the remaining parameter are the same as those which will be used for simulations later.

this filter is show in figure 5.1. For an off-center/on-surround receptive field, the center Gaussian is negative, while the surround is positive. Mathematically, the model

can be expressed by [93]

$$D_s(x,y) = \pm\big[\frac{A_c}{2\pi\sigma_c^2}e^{-(x^2+y^2)/2\sigma_c^2} - \frac{A_s}{2\pi\sigma_s^2}e^{-(x^2+y^2)/2\sigma_s^2}\big] \tag{5.6}$$

where A_c and A_s are scaling factors for the center and surround component, and σ_c and σ_s defines the spread of the center and surround. The $\pm$ sign indicates whether the receptive field is on-center $(+)$ or off-center $(-)$.

5.3 Responses of Spatial Filters

In the following, the response of the single Gaussian and an on-center Difference-of-Gaussians spatial filter to two different stimuli is calculated from equation (5.2). For the off-center filter, the response is just the inverse of this result. The stimuli are full-field drifting gratings and flashing spots of uniform luminosity. In addition, the response of the on-center DoG filter to patches of drifting grating will be described. More specifically, in subsection 5.3.1, the spatial frequency dependence of the response to drifting gratings is shown. Subsection 5.3.2 describes how the response to flashing spots varies as a function of spot diameter. The diameter dependence of the response to drifting patch gratings is considerably more complex, but will be briefly discussed in section 5.3.3. The three stimuli will later be used as visual input to drive the Allen model.

5.3.1 Response to full-field drifting gratings

Using horizontally moving gratings as visual input, we have

$$S(x,y,t-\tau) = A\cos\left(kx - \omega(t-\tau)\right) \tag{5.7}$$

where A is the contrast, k is the wavenumber $k = 2\pi\nu$ and ω is the angular frequency $w = 2\pi f$. The spatial and temporal frequencies, ν and f, are measured in cycles per degree (cpd) and Hz respectively. The contrast A was set to one and omitted in the following. By using the two-dimensional Gaussian function in equation (5.4) as a spatial filter, we can integrate over the spatial component in equation (5.2). For a full field drifting grating, the spatial boundaries in this integral extend from minus infinity to infinity. The resulting spatial frequency dependence of the response is thus

$$L_s(k) = e^{-k^2\sigma^2/2}. \tag{5.8}$$

where σ defines the spread of the Gaussian.

To model center-surround rather than center-only receptive fields, the spatial filter model can be modified from a simple Gaussian to the Difference-of-Gaussians [93] in equation (5.6). Using this and equation (5.7), the result is a filter response with the spatial frequency dependence of

$$L_s(k) = A_c e^{-k^2\sigma_c^2/2} - A_s e^{-k^2\sigma_s^2/2}. \tag{5.9}$$

The derivations of equation (5.8) and (5.9) are presented in appendix A. The two results are illustrated in figure 5.2. The single Gaussian produces a pure low-pass response, while the double Gaussian gives rise to the band-pass tuning curve associated with linear surround suppression. The total response of a DoG filter is given by

$$L(t) = |\Lambda(\omega)|\left[A_c e^{-k^2\sigma_c^2/2} - A_s e^{-k^2\sigma_s^2/2}\right]\cos\left(\omega t - \theta(\omega)\right). \tag{5.10}$$

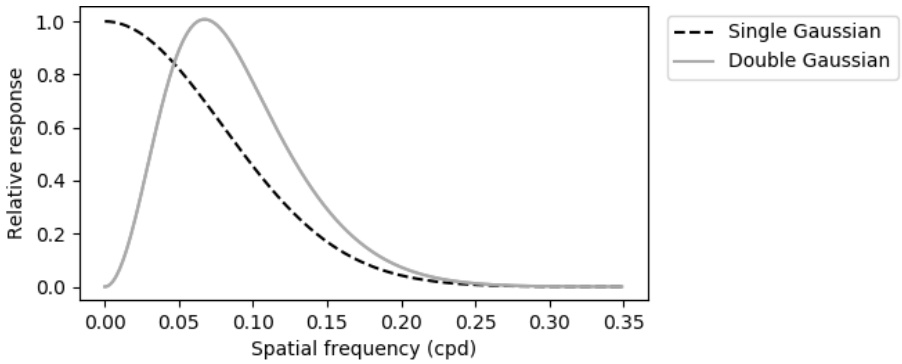

Figure 5.2: The analytical solution to the response amplitude of Gaussian spatial filters as a function of spatial frequency of full field drifting gratings. The dashed black line show the response given by equation (5.8) with $\sigma = 2$, for single 2D Gaussian spatial kernels. The solid grey line shows the DoG response, given by equation (5.9), with parameters $A_c = A_s = 1.73$, $\sigma_c = 2$ and $\sigma_s = 2.45\sigma_c$. The value of σ and σ_c refer to a representative cell in the LGN model, while the remaining DoG parameter are the same as those which will be used for simulations later.

where $|\Lambda(\omega)|$ and $\theta(\omega)$ is determined by the temporal filter (see appendix A). Hence, the response is an oscillation as a function of time, where equation (5.9) gives the spatial dependence of the amplitude. The curves in figure 5.2 therefore show the tuning of the mean firing rate amplitude, rather than just the mean firing rate.

5.3.2 Response to flashing spots

The spatial filter response to an isoluminant spot can be expressed mathematically as a function of spot diameter. For a circular, isoluminant stimuli of a fixed radius, integral 5.2 becomes the integral over a circular disc. The stimulus is now a constant in time, so that the spatial response is defined by the integral of $D_s(r)$ over the circle. If a single Gaussian is substituted for $D_s(r)$, one can calculate the diameter dependence of the response to be

$$L_s(a) \propto (1 - e^{-a^2/2\sigma_c^2}) \tag{5.11}$$

where a is the radius of the circle. If the double Gaussian in equation (5.6) is used as the kernel, the result is

$$L_s(a) \propto \left[A_c(1 - e^{-a^2/2\sigma_c^2}) - A_s(1 - e^{-a^2/2\sigma_s^2})\right]. \tag{5.12}$$

See appendix A for more details on the derivation of equations 5.11 and 5.12. Figure 5.3 illustrates the solutions given by equation (5.11) and (5.12). The single filter creates a response which reaches a maximum when the stimulus covers the whole spatial kernel, and does not decrease for larger stimuli. For the double Gaussian on the other hand, the response has a peak when the stimuli covers the whole center Gaussian, but stimuli larger than this will inhibit the activity. That is, the filter exhibits surround suppression of uniform stimuli.

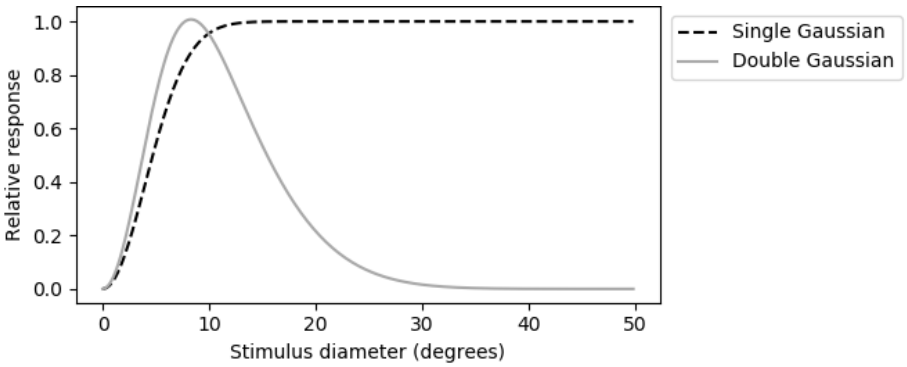

Figure 5.3: The analytical solution to the response of a Gaussian filter, given by 5.12, to an isoluminant circle, plotted as a function of stimulus diameter. A single 2D Gaussian spatial kernel produces the response illustrated by the dashed black line, given by equation (5.11), where $\sigma = 2$. The DoG response is shown by the solid grey line, given by equation (5.12). The parameters for the DoG are $A_c = A_s = 1.73$, $\sigma_c = 2$ and $\sigma_s = 2.45\sigma_c$. The value of σ and σ_c refer to a representative cell in the LGN model, while the remaining DoG parameter are the same as those which will be used for simulations later.

5.3.3 Response to patch gratings

In section 5.3.2 it is explained how the DoG filter generates surround inhibition for increasing sizes of uniform, bright stimuli. For increasing sizes of circular patches of drifting grating, the response is more complex. In this case, the activity predicted by the DoG filter will not have the same diameter dependence as for uniform stimuli. In fact, for certain combinations of spatial frequencies and filter parameters, the surround effect can turn in to surround enhancement [95]. Similarly for flashing spots and full field drifting gratings, the response of the spatial kernel in integral 5.2 can be evaluated for circular patch gratings. This response can be expressed as a function of the wavevector, k and radius, a. The final result, as derived by [95], is given by

$$L_s(a,k) = A_c X(\frac{\sigma_c}{\sqrt{2}a}, \sqrt{2}\sigma_c k) - A_s X(\frac{\sigma_s}{\sqrt{2}a}, \sqrt{2}\sigma_s k) \tag{5.13}$$

with

$$X(y,z) = e^{-z^2/4}\frac{1}{4y^2}\sum_{n=0}^{\infty}\frac{1}{n!}(\frac{z}{2})^{2n}\,{}_1F_1(n+1;2;-1/4y^2) \tag{5.14}$$

where $z = 2ayk$ and ${}_1F_1$ is a series called "confluent hypergeometric function" [95]. Consequently, a large variety of responses may be generated with the DoG filter from circular patch grating stimuli. Some example curves are shown in figure 5.4. The size tuning of a cell in this case depends on the combination of the grating's spatial frequency and the DoG parameters. Two distinct cases leading to surround enhancement and surround suppression can be seen in figure 5.4. In the top of figure 5.4, the DoG has parameters so that the ON-center Gaussian overlaps the light region at the same time as the OFF-surround overlaps (mostly) dark areas. Since ON-regions respond to bright stimuli, and OFF-regions respond to dark stimuli, this combination of center and surround causes a higher activity than the response of the DoG to uniform input of the same diameter. The grating therefore has the preferred spatial

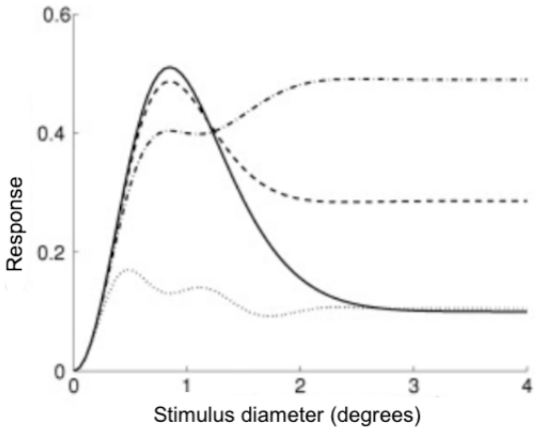

Figure 5.4: Response of the DoG filter model versus diameter of a patch of drifting grating for different values of k. The dashed lines show surround suppression ($k = 2$) and surround enhancement ($k = 4.4$). The grey line corresponds to $k = 10$. The solid black line illustrates a case where $k = 0$, which is an isoluminant spot. Note that the DoG parameters are different to those in figure 5.1, 5.2 and 5.3. Adapted from Figure 3 in [95].

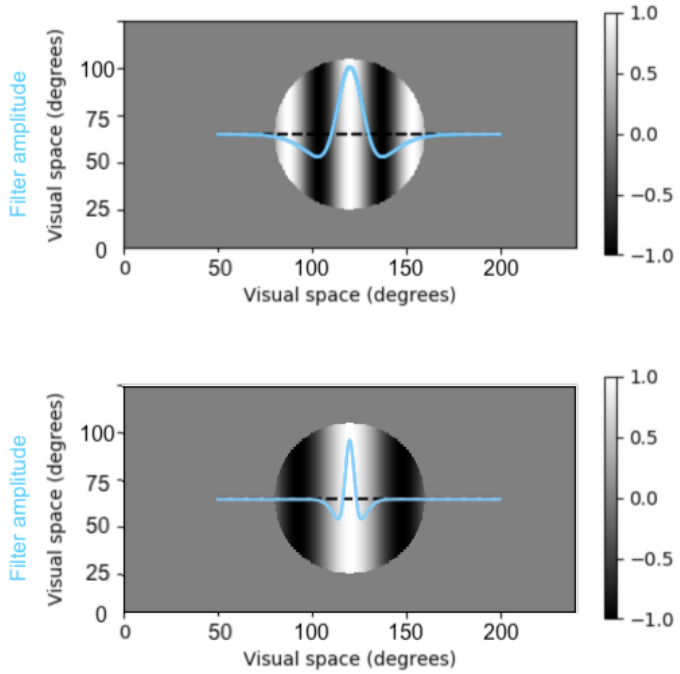

Figure 5.5: A combination of a DoG filter and patch grating which would lead to surround enhancement (top) and surround suppression (bottom).

frequency for this DoG filter. The figure also illustrates why activity continues to increase with stimuli size until it covers both the center and surround, and thereby

giving surround enhancement. The bottom figure shows an example where the DoG center and surround is small compared to the grating spatial frequency. In this case, the response of the DoG to the grating approaches the response to uniform stimuli, and surround suppression can be observed. Similarly, Einevoll and Plesser [95] predicted more surround suppression for low spatial frequencies compared to high, as demonstrated in figure 5.4.

5.4 The Gabor Filter

The single Gaussian and Difference-of-Gaussians are not the only functions used to describe receptive fields. A different, but related, model is the Gabor function. The function was described by Dennis Gabor in 1946 [94]. A few decades later, it was extended to two dimensions and used to describe receptive fields of cells in the visual cortex [97–99]. The model is able to describe the spatial frequency tuning of the cells, as well as their preference for orientation and localised regions in space [99, 100]. The spatial filter, corresponding to the real part of the complex Gabor function, is a product of a sinusoidal function and a Gaussian envelope [50],

$$D_s(x,y) = \frac{1}{2\pi\sigma_x\sigma_y} e^{-(x^2/2\sigma_x^2)-(y^2/2\sigma_y^2)} \cos\left(kx-\phi\right) \tag{5.15}$$

where σ_x and σ_y define the size of the receptive field and k and ϕ are the preferred spatial frequency and phase of the filter.

While the circular center-surround receptive fields in LGN and retina cells are well approximated by the Difference-of-Gaussian, the Gabor function is often the preferred model for visual cortex cells [101]. Comparisons between the mathematical model and experiments show that the behaviour of cells in the visual cortex agree well with the Gabor filter description [98, 99, 102]. That said, while some refer to the Gabor function as the most robust model for V1 receptive fields [103], others argue that in comparisons with other filters, alternate models have provided better fits to V1 experimental data, including the Difference-of-Gaussians [104] and the related Gaussian derivative model [101, 105].

Chapter 6

The Allen Model

The following chapter is primarily based on the pre-print publication on the Allen model by Billeh, et al. [12], as well as on the Allen model scripts and data, available from [106].

6.1 The Allen Network Model

This thesis explores the behaviour of a mouse primary visual cortex (V1) model. The model was developed by the Allen Institute for Brain Science, as a step towards gaining insight into the human brain. It was constructed on the basis of large-scale literature reviews and experimental measurements. The V1 network is made up of point neurons in the form of GLIF cells with after-spike currents [55]. For a general description of this type of model and its dynamics, see section 3.1.2. The GLIF network is one of two analogous V1 models developed by the Allen Institute. The other uses a biophysically detailed representation of the neuron. Otherwise, the two models are identical, having the same connections, coordinates and distribution of cell classes. The reason for choosing to work with the GLIF model is primarily the computational efficiency ($\sim$ 8000 times faster than its counterpart, according to the Allen Institute [12]). This enables simulations to be run outside a supercomputer. Additionally, the two models perform very similarly, both qualitatively and quantitatively [12].

The Allen V1 model consists of 230 924 cells, compared to the 360 000 nerve cells in real mice [107]. The cells are distributed uniformly in the physical space, so that the network forms a cylinder with a 845 µm radius. Similar to real mice, the network is divided into the cortical layers LI, LII/III, LIV, LV and LVI, as illustrated in figure 6.1. Cell coordinates and connections are set by well-defined stochastic processes. Individual neurons are randomly drawn from models of 17 different classes. Each class and model is based of real neuron types, fit to electrophysiological data from the Allen Institute's own database. The neurons are divided into classes based on their gene expression related to electrophysiological responses [108].

As illustrated in figure 6.1, each V1 layer consist of a specific set of neuron classes. Neurons in LI are of a serotonin-expressing (Htr3a) inhibitory class. The remaining layers contain both excitatory and inhibitory neurons. The inhibitory neurons are separated into the classes of Htr3a, somatostatin-expressing (Sst) and parvalbumin-expressing (Pvalb) cells, which have been found to make up most of the inhibitory cells in V1 of real mice [109]. The excitatory class is separated into one E-class per layer. The proportions of excitatory and inhibitory neurons are 85% and 15% respectively.

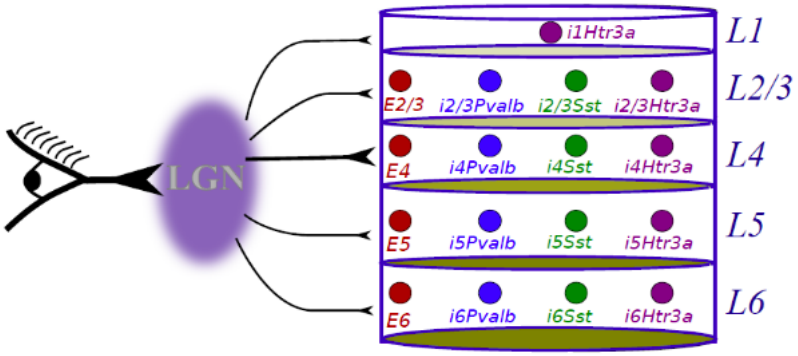

Figure 6.1: A representation of the Allen LGN and V1 model. The LGN receives visual input and sends its output to V1. The illustration shows the cell classes and the six (practically five) layers of the V1. The classes are further subdivided into different models, based on slice electrophysiology measurements performed by the Allen Institute. Reprinted from Figure 1 in [12].

The V1 model requires input in the form of neuronal spike trains. For this purpose, the Allen Institute developed a representation of the LGN in the form of spatio-temporal filters units. The LGN module takes arbitrary visual stimuli as input and generates spike trains which are relayed to V1. The V1 neurons which receive LGN input are E - and Pvalb cells in LII/III-LVI, as well as Htr3a cells in LI. The Sst and Htr3a neurons in LII/III, on the other hand, receive no such input. The probability that a specific V1 cell receives LGN input further depends on the neuron class and layer. In accordance with experiments [110], more LGN axons transmit to LIV than any other layes.

The connections between the LGN and V1 were established stochastically, with the aim of generating orientation and direction selectivity in V1. Connection probabilities depend on the type of cell and their spatial position. For each V1 cell, two elliptical regions within its receptive field were defined, aligned along the axis of the neuron's preferred orientation. Connections were then established between the V1 cell and a set of LGN cells with receptive fields within these regions. From one region, LGN cells with sustained (slow) kinetics were sampled, while for the second region, LGN cells with transient (fast) responses were selected. This creates different V1 responses to stimuli of opposing directions, due to different timings of LGN input in the two cases. Figure 6.2 illustrates how connections from LGN cells to a single V1 cell were selected, and how this gives rise to direction selectivity in the V1 cell.

Connection probability and synaptic strength for recurrent connections within V1 were set depending on distance, cell type and similarity in preferred orientation, phase and direction of motion of the visual stimulus. Rules for synaptic connections are based on experimental literature, as well as measurements performed at the Allen institute. These rules depend on the classes of the connected pairs, and typically, cells with similar stimuli preferences have stronger synaptic weights and higher probabilities for establishing connections.

The distance-dependence of connections follows the Gaussian function in equation (3.12), with parameters σ_r and A. The standard deviation, σ_r, takes different values between 85 and 120 µm, depending on cell types (E, Sst, Pvalb, Htr3a; independent of layers). The connection probability therefore decreases with distance measured in

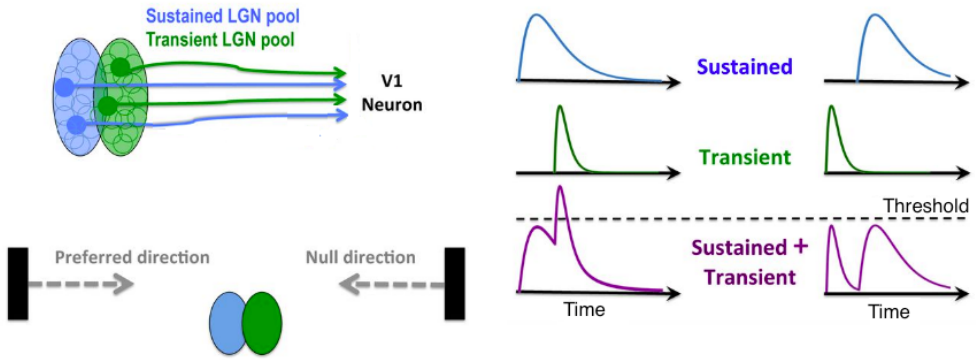

Figure 6.2: The connections between LGN and V1 cells are selected to allow for orientation and direction selectivity in the V1 cells. For stimuli in the preferred direction, the slow and fast responses would overlap in time, causing the membrane potential in the V1 neuron to exceed the firing threshold. For the opposing direction, however, the input does not overlap, resulting in a weak or no V1 response. Adapted from Figure 2 in [12].

the cortical plane. Distance in cortical depth was not considered. The scaling factor A, however, depends on both cell types and layers, were cells in the same layer are generally more likely to be connected than cells from different layers. The GLIF model uses current-based synapses, as described in section 3.3, with dynamics governed by equation (3.15).

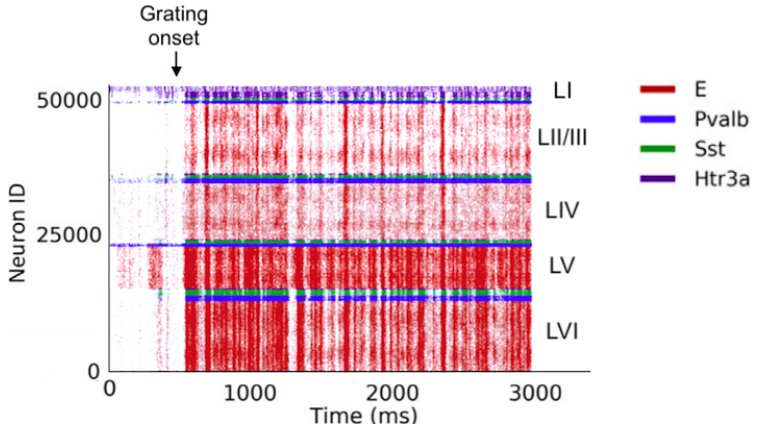

Figure 6.3: Example raster plot of the output from a V1 simulation. The model produces output in the form of neuronal spike trains. Each spike is represented by its firing time and each neuron by its neuron ID. Spike trains for a proportion of the V1 neurons is shown. For this example, the visual stimulus was a drifting grating of 2 Hz and 0.04 cpd, preceded by a grey screen of 500 ms duration. Adapted from Figure 5a in [12].

In addition to receiving input from the LGN and other V1 neurons, V1 neurons receive "background input" in the form of a Poisson spike train, to approximate the input from other brain areas. The Poisson spike train has a constant frequency of 1000 Hz and is relayed to all V1 cells. The connection strengths from the background

input were tuned so that each neuron generated the right amount of spontaneous activity according to experiments.

The receptive field centers of V1 neurons are defined via a mapping between the visual space and the physical cell locations. Each location in the cortical plane of the V1, spanned by the x and z-axes, correspond to coordinates in the visual input space. The center of the two spaces are mapped onto each other, and from here, one millimeter in the cortex correspond to 40 °/mm in elevation and 70 °/mm in azimuth in the visual space.

The output from the V1 model takes the form of neuronal spike trains. An example raster plot of results from a single simulation is illustrated in figure 6.3. The figure gives an overview of the different layers in the model and shows the relative proportions of the various cell types.

6.2 The Allen LGN Model

The Allen V1 model receive spike train inputs from the LGN. These spike trains are created in an LGN module consisting of 17 400 units. In comparison, the mouse LGN contains approximately 18 000 neurons [107]. LGN cells are modelled as Linear-Nonlinear Poisson neurons (see section 3.2.1). Each cell takes visual stimuli of dimensions 240°×120° as input, on which it performs a linear transformation. This linear transformation takes the form of a separable spatio-temporal filter (see section 5.1). The output is then passed through a nonlinear rectifier, which includes adding a spontaneous firing rate and applying a threshold to avoid negative firing rates. The rectifying function is given by equation (3.10). The result is a time series of firing rates for each unit [12]. After calculating the firing rates, spike trains are generated from the rates through a Poisson process. See figure 6.4 for an illustration of how the visual input is transformed into spike trains.

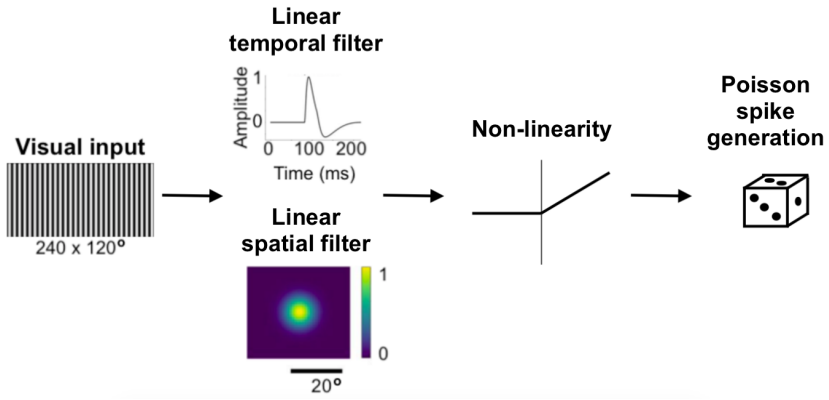

Figure 6.4: The Allen LGN model is a type of Linear-Nonlinear Poisson model. Linear spatial and temporal filters are first applied to the input. The result is then passed through a non-linearity, giving the instantaneous firing rate of the neuron. A Poisson process is then used to generate stochastic spike trains based on the firing rate. The temporal and spatial filter illustrations are examples for the sON-TF8 cell type and are adapted from Figure 2 of the Allen documentation [12].

Four classes of cells were developed, based on experimental responses measured by [16] and [52]. These were one sustained ON class, a sustained OFF class, a transient OFF class and a group of ON/OFF cells. These classes are further divided into

subclasses (TF1, TF2, TF4, TF8 and TF15), according to their preferred temporal frequency of stimuli. Temporal selectivity depends on the shape of the temporal filter. This filter is defined in equation (5.5) and illustrated in figure 5.5, where parameters are different for each cell. The sustained and transient units have slow or fast kinetics respectively, depending on the width of the temporal filter.

The coordinates of the LGN cells refer to points in visual space, rather than explicitly modelling their physical location. The cells are distributed over a $240° \times 120°$ visual plane, with coordinates specifying the center of their spatial filters. The ON or OFF type cells have a single circular filters, while the ON/OFF cells consist of two circular subunits: one ON region, and one OFF region. The majority of cells are either a simple ON or OFF unit. The spatial filter for a single ON or OFF cell takes the form of the two-dimensional Gaussian kernel in equation (5.4). The spatial field sizes vary between cells, as defined by $\sigma \in [0.67°, 3.31°]$[1]. The ON-cells have positive Gaussians, while the OFF-cells have negative. For the ON/OFF cells, the spatial filter consist of two 2D Gaussians with opposite polarity, separated by a few degrees [111]. The Gaussian of one subunit has a larger amplitude and dominates the other [111].

[1]The value of σ is defined in the script `lgn_functions.py` [111] as "spatial size" divided by 3. Here the values have been calculated using the "spatial size" parameter for each cell, extracted from `lgn_full_col_cells_3.csv` [112].

Chapter 7

Experimental Findings

In this chapter, experimentally obtained spatial and temporal tuning properties of mice LGN and V1 are presented. The aim is to establish the phenomena observed in the mouse brain, in order to evaluate the performance of the Allen model. Importantly, experimentally observed spatial properties have been used as a basis for modifying the Allen model. The properties looked at are size tuning, spatial frequency tuning and temporal frequency tuning.

As described in section 4.1.1, different mechanisms are likely involved in size suppression of different stimuli. For example, a linear receptive field with on-off antagonisms can explain suppression in response to uniform stimuli. However, size suppression which can only be explained by nonlinear mechanisms have also been found in mammals (see section 4.1.1). This type of suppression is better measured with patch gratings, as the suppressive effect of linear mechanisms is usually much smaller for this type of stimuli [63, 66] (see section 5.3.3 for a mathematical explanation). When looking at experiments in mice, it is therefore preferable to include size tuning experiments with both uniform and patch grating stimuli. However, no studies on mouse LGN with patch gratings have been found.

Nevertheless, experiments on LGN and V1 in mice for a variety of stimuli are described in this chapter. The studies on both LGN and V1 investigated spatial and temporal frequency tuning using drifting gratings, as well as size tuning using flashing spots. In addition, studies on V1 investigated size tuning with patch gratings. This chapter primarily describes the studies which will later be referenced when comparing the Allen model to real mice. Each study is briefly summarised with respect to the property of interest. For a general overview, each subsection is concluded with a short summary and tables conveying the most important findings. In addition, each section includes a discussion of possible effects of anaesthesia on the results.

7.1 Spatial Tuning in the Mouse

7.1.1 LGN

Grubb and Thompson [113]

Grubb and Thompson [113] measured the response of single neurons in LGN of adult anaesthetised mice. Full field drifting gratings were used as visual stimuli to investigate spatial frequency preference. Keeping the contrast at 70 % and the temporal frequency at 1 Hz, gratings of different spatial frequencies were presented. Difference of Gaussian curves were successfully fitted to 92 cells. The majority of cells showed

band-pass spatial frequency tuning, with a preferred frequency of ~ 0.03 cpd (cycles per degree). Additionally, receptive fields were identified in 133 out of 199 cells by using stimuli of various sized black or white flashing squares. The remaining cells did not have clearly defined regions which reacted reliably to visual stimuli. The receptive fields were circularly symmetric, with a large proportion showing center-surround antagonisms. [113]

Piscopo et al. [52]

Piscopo et al. [52] is one of the experimental studies on which the Allen Institute based their LGN model [12]. The study recorded various response properties in the LGN of 16 adult anesthesised mice. Neurons were classified into six groups based on their properties. The first three groups consisted of sustained ON (sON), sustained off (sOFF) and transient OFF (tOFF) cells, from a total 134 out of the 214 responsive neurons. These cells were found to have center-surround receptive fields. In addition, other cell types with more complex properties were also discovered. Among these were orientation and direction selective neurons, a group of slowly responding neurons and a group of cells that were suppressed by nearly all types of stimuli. Flashing, uniformly black or white spots of increasing diameters were used to measure surround suppression of the sON, sOFF and tOFF cells. Close to all neurons in the three center-surround groups (sON, sOFF and tOFF) demonstrated strong surround suppression by a decrease in firing rate for diameters larger than preferred size. In the same study, they also measured the response of these cells to drifting gratings of different spatial frequencies. The responses showed band-pass tuning with peak activity at 0.05, 0.03 and 0.02 cpd for the sON, sOFF and tOFF cells respectively. [52]

Tang et al. [28]

Tang et al. [28] recorded neurons in the LGN of aneasthetised, adult mice. After excluding non-responsive neurons, receptive fields were mapped in 185 cells, using movies of stochastic noise with varying contrasts. All recorded cells displayed center-surround antagonistic receptive fields. Full-field drifting gratings were then presented as stimuli for spatial frequencies of 0.02 to 0.96 cpd. The spatial frequency tuning of each cell was then successfully fitted to a DoG function for 92/185 cells. Preferred spatial frequencies were taken as the peak calculated from this DoG function. 51 % of the cells preferred spatial frequencies below 0.02 cpd, and were hence classified as showing low pass tuning. The remaining 49% showed band-pass tuning, with a median preferred spatial frequency of 0.035. [28]

Durand et al. [16]

Durand et al. [16] studied properties including the optimal spatial frequencies of cells in LGN. While previous studies on mice LGN have focused on anaesthetised animals, this study included measurements on adult mice in the awake and anaesthetised state. The number of LGN cells recorded from were 456 and 164 in the awake and anaesthetised state respectively. Along with the study by Piscopo et al. [52], this experiment was used as a benchmark by the Allen Institute when developing LGN filters and tuning the cell parameters [12]. Spatial frequency tuning curves were constructed for single cells, and the proportion of cells firing optimally over a range of frequencies was measured. The distribution of peak frequencies was found to follow a wide curve, with preferred spatial frequencies in the range 0.02 - 0.64 cpd. In awake mice LGN, the median optimal spatial frequency was 0.08 cpd. [16]

study	stimuli	state	amount of size suppression
Piscopo et al. [52]	flashing spots	anaesthetised	nearly all sON, sOFF & tOFF cells
Tschetter et al. [83]	flashing spots	awake	mean SI: 0.53
Grubb & Thompson [113]	flashing spots	anaesthetised	large proportion

Table 7.1: Overview of experiments measuring size suppression in the mouse LGN. The column "amount of size suppression" is the degree of surround suppression found, either given with suppression index (SI) or as the amount of tested cells that show surround suppression. For the study by Tschetter et al., only results for for adult mice are included.

Tschetter et al. [83]

Tschetter et al. [83] compared surround effects in the LGN of mice of different age groups. Measurements were made on awake mice. Circular black or white flashing spots of different diameters were presented. The study was conducted by performing measurements on mice across groups depending on age. 228 cells were recorded in a group of 17 adult mice (45 to 60 days old). Groups of developing mice of 14-16, 17-21 and 22-24 days old consisted of 7, 22 and 7 mice respectively. The numbers of neurons recorded were respectively 102, 298 and 115. The size tuning curve showed qualitative differences between developing mice and adult mice. The youngest group showed a strong increase in firing rate with diameters up to full field luminance. Adult mice, however, showed size suppression, with increased response up to a preferred diameter, and then a decrease in firing rate. The preferred spot diameter in adult mice was 8.6°, and the suppression index was 0.53. The largest suppression index, 0.61, was found in mice aged 22-24 days. (Recall section 4.1.2 for a definition of suppression index.) Flashing spots at different locations in visual space were used to map receptive fields, which often revealed center-surround structures. The spatial frequency tuning curves were also measured for the different age groups. The results were consistent with the age dependence described above. There was an increase in preferred frequency from 0.02 cpd in the youngest group to a peak at 0.12 cpd in the group aged 22-24 days. For adult mice, the preferred spatial frequency was 0.09 cpd. [83]

Summary of findings

In this section, the findings of five studies on the mouse LGN are summarised. The main relevant results are shown in the tables 7.1 and 7.2 for size tuning and spatial frequency analysis respectively. The studies on size tuning all found cells with surround suppression [52, 83, 113]. The exact number of cells showing size suppression was not specified for all studies, but the property was found in a large proportion of cells in the three relevant studies [52, 83, 113]. Consistent with this, the same exper-

study	stimuli	state	pref sf (cpd)
Piscopo et al. [52]	full-field grating	anaesthetised	means: 0.05 (sON), 0.03 (sOFF), 0.02 (tOFF)
Durand et al. [16]	full-field grating	awake & anaesthetised	median: 0.08 (awake)[a]
Tschetter et al. [83]	full-field grating	awake	mean: 0.09
Grubb & Thompson [113]	full-field grating	anaesthetised	most cells: $\sim$ 0.03
Tang et al. [28]	full-field grating	anaesthetised	median: 51% of cells: <0.02 49% of cells: 0.035

Table 7.2: Overview of experiments studying spatial frequency tuning in the mouse LGN, using full-field drifting grating as stimuli. For the study by Tschetter et al. [83], only results for adult mice are included.

[a]Durand et al. [16] performed measurements in awake and anaesthetised mice, but the median preferred frequency was only reported for awake mice. However, anesthesia was said not to affect the spatial frequency preference [16].

iments found that a large number of cells had receptive fields with center-surround antagonisms [52,83,113]. Also compatible with this, a large proportion of cells showed band-pass spatial frequency tuning [16,28,52,83,113]. Different preferences for spatial frequencies were found in the five studies. In the three studies looking at anaesthetised mice, a large portion of cells had peak frequencies around 0.03 cpd [28, 52, 113]. In the two studies recording neurons in awake mice, the LGN neurons had preferred spatial frequencies of 0.08 cpd (median) [16] and 0.09 cpd (mean) [83]. Still, Durand et al. [16] found no statistically significant difference between the mean peak frequencies of awake and anaesthetised mice. Durand et al. [16] suggests that the difference between the optimal spatial frequencies can be ascribed to the use of different temporal frequencies. In previous studies, temporal frequencies of 1 Hz [113] and 2 and 8 Hz [52] were used to find the peak spatial frequencies, while Durand et al. [16] used 1, 2, 4, 8 and 15 Hz. Tschetter et al. [83] found that surround suppression is age-dependent, mostly manifesting in mice of around 22 days and older. Mice classified as adult[1] (45 to 60 days) showed significant surround suppression, whilst very young mice did not [83]. In conclusion, the studies imply that a large amount of LGN cells demonstrates surround suppression in response to uniform stimuli, consistent with the band-pass spatial frequency tuning results.

7.1.2 V1

Schuett et al. [48]

Schuett et al. [48] mapped neurons in several areas of the mouse cortex, including V1. Measurements were made for mice in the anaesthetised state. As stimuli they used square patches of drifting square wave gratings at different positions in the visual field. This stimuli induced increasing activity when presented in the center of a neuron's receptive field. An activity below the spontaneous firing rate was considered as surround inhibition. Out of 22 recorded cells, 17 displayed surround inhibition. [48]

Durand et al. [16]

In addition to investigating cells in LGN, Durand et al. [16] studied how cells in V1 responded to different stimuli. For awake mice, 232 V1 cells were recorded, while 80 V1 cells were recorded in aneasthetised mice. As when measuring LGN cells, they used drifting gratings of different spatial and temporal frequencies, and thus determined the optimal spatial frequencies for the V1 cells. Compared to in the LGN, the distribution of peak spatial frequencies where narrower, but still ranged from 0.02 to 0.32 cpd. For awake mice, the median optimal spatial frequency was 0.05 cpd in the V1, compared to 0.08 in the LGN, but this difference was not statistically significant. The optimal frequency for anaesthetised mice was not specified, but was reported not to be statistically different from that in awake mice. [16]

Adesnik et al. [115]

Adesnik et al. [115] recorded neuronal activity in V1 layer 2/3 of 1-3 months old awake mice. Mice were presented with drifting patch gratings of 0.04 cpd spatial frequency as visual stimuli. Patch sizes were varied from 8° to 96 ° in diameter. Out of recordings from 53 random single neurons, 33 neurons showed significant surround suppression,

[1]The definition of adulthood in mice varies across literature, and has been reported to start anywhere from 6 to 20 weeks old [114].

with a preferred diameter of 22° ± 2°. The mean suppression index (as defined in section 4.1.2) for single cell recordings was 0.9. In addition, targeted recordings of somatostatin expressing neurons (Sst) and parvalbumin expressing neurons (Pvalb) were performed. The mean suppression indices for Pvalb and Sst neurons were 0.46 and 0.09 respectively. While 6 out of 11 cells showed significant surround suppression in Pvalb neurons, the suppression was not significant in any of the 8 recorded Sst cells. This indicates that size tuning depends on the type of neuron.

Additionally, the study looked into recurrent connections in V1 and how these may give rise to some of the suppression observed. In particular, they found that the activity of Sst neurons increased consistently as a function of stimuli diameter, for diameters up to 900 µm. As Sst neurons receive little or no input from LGN neurons [12, 115], the increase in firing rate was taken as being due to excitatory input from LII/III pyramidal cells, which were found to provide the main input to Sst neurons. To further confirm this, direct light stimulation of slices of V1 were used. It was then found that creating two vertical cuts through LII/III prevented the firing rate in Sst neurons from increasing for stimuli extending past these cuts. The findings suggest the existence of far reaching horizontal connections within V1, which gives rise to a preference for large stimuli in Sst neurons. The increased inhibition by Sst neurons for large stimuli could in turn cause surround suppression in the remaining cell types. [115]

Self et al. [64]

Self et al. [64] performed measurements on different layers of V1 in 5-8 weeks old anaesthetised mice. Drifting gratings of various spatial frequencies, temporal frequencies and orientation were presented, and the most common preferred parameters for the recorded group of cells were found. Surround suppression was then measured using these parameters, for drifting gratings of various sizes. 106 neurons were measured, and significant differences in the strength of surround suppression between V1 layers were found. Single neuron recordings gave a median suppression index (as defined in section 4.1.2) of 0.37 in LIV. Multi cell recordings in deep layers (-500 to -100 µm cortical depth), LIV (0 to 100 µm) and superficial layers (200 to 400 µm) gave median suppression indices of 0.06, 0.25 and 0.18 respectively. Median preferred stimulus sizes ranged from 17.5° in LIV to 54° in LVI. [64]

Van den Bergh et al. [116]

Van den Bergh et al. [116] recorded 69 neurons in V1 of eight adult anaesthetised mice. In order to determine size tuning properties of the V1 neurons, circular patches of drifting gratings of diameters in the range 1° - 70° were presented. The spatial frequency, temporal frequency, contrast and orientation of the grating were set to optimal values for each recorded neuron. Pronounced surround suppression was found in most of the recorded V1 neurons, although about 40 % of units did not show surround suppression. The mean suppression index was 0.24. The median preferred diameter was 28.9°. In addition to size tuning, the spatial frequency tuning of V1 cells was measured. The mean preferred spatial frequency was 0.04 cpd. [116]

Vaiceliunaite et al. [82]

Vaiceliunaite et al. [82] used circular patch gratings of increasing sizes to measure the size tuning of V1 cells of the adult mice. The spatial and temporal frequencies were set to the average optimal values, and diameters ranged from 4° to 67°. Measurements

study	stimuli	state	amount of size suppression
Schuett et al. [48]	patch gratings	anaesthetised	17/22
Adesnik et al. [115]	patch gratings	awake	any type: 33/53, mean SI: 0.9 Pvalb: 6/11, mean SI: 0.46 Sst: 0/8, mean SI: 0.09[a].
Self et al. [64]	patch gratings	anaesthetised	mean SI, single: 0.37 (LIV) mean SI, multi rec: 0.06 (deep layers) 0.25 (LIV) 0.18 (superficial layers)
Van den Bergh et al. [116]	patch gratings	anaesthetised	$\sim$ 60 %
Vaiceliunaite et al. [82]	patch gratings	awake & anaesthetised	65% (awake) 47% (anaesthetised)
Dräger [117]	flashing spots & bars	anaesthetised	large proportion

Table 7.3: Experimental studies on size suppression in the mouse primary visual cortex. The column "amount of size suppression" is the degree of surround suppression found, given with suppression index (SI) or as the amount of tested cells that showed significant suppression.

[a] Although mean SI was 0.09, the observed suppression was not statistically significant in any Sst cells [115].

study	stimuli	state	pref sf (cpd)
Durand et al. [16]	full-field gratings	awake & anaesthetised	median: 0.05 (awake)[a]
Van den Bergh et al. [116]	full-field gratings	anaesthetised	mean: 0.04
Marshel et al. [118]	full-field gratings	anaesthetised	mean: 0.045

Table 7.4: Experimental studies on spatial frequency tuning in the mouse primary visual cortex.

[a]Only the median preferred spatial frequency for cells in awake mice, not anaesthetised, was specified by Durand et al. [16], though optimal frequency was reported not to be affected by state.

were performed in both awake (68 cells) and anaesthetised mice (163 cells). Most cells (65%) in the awake mouse showed size tuning consistent with surround inhibition, while 47% showed suppression in the anaesthetised mouse. [82]

Marshel et al. [118]

Marshel et al. [118] mapped spatial frequency tuning in multiple cortical areas of mice, including V1. The 28 mice were anaesthetised during the recordings. The stimuli consisted of drifting gratings of five different spatial frequencies, from 0.01 to 0.16 cpd. A total of 728 neurons were responsive and reliable for the V1 spatial frequency analysis and hence included in the analysis. The mean preferred spatial frequency was 0.045 cpd. There was a large variation between peak spatial frequencies, with most cells showing band-pass tuning, a large fraction showing high-pass tuning and fewest cells showing low-pass tuning. [118]

Dräger [117]

Dräger [117] measured the receptive field structure of cells in the primary visual cortex of mice. Recordings were made from 400 single cells and some cell clusters in 4 to 8 months old anaesthetised mice. Cells were categorised into subgroups, based on their responses to flashes of light and various moving stimuli. One group preferred moving, elongated stimuli of a certain orientation (45% of cells). The other group had circularly symmetric receptive fields with no orientation preference (55%). Stationary flashes of light were used to map the cells' receptive fields. Out of the circularly symmetric cells, over half had a center-surround structure with an on-off antagonism, where the surround region inhibited the response of the center. The others had no opposing surround. Out of the orientation selective cells, the simple cells[2] (17% of all cells) had elongated, antagonistic on-off areas. Most common was an elliptic on-region in between two off-regions, so that responses to flashing stimuli was inhibited when stimuli extended into the off-regions. The remaining cells preferring oriented stimuli, termed complex and hypercomplex, had apparently uniform receptive fields.

[2]Recall section 2.2.2 for the definition of simple and complex cells.

All types of cells responded in general weakly or not at all to full-field changes in luminosity. [117]

Summary of findings

The studies on mouse V1 found that a large proportion of cells show band-pass spatial frequency tuning. The preferred frequencies found were 0.04 cpd (mean) [116], 0.045 cpd (mean) [118] and 0.05 cpd (median) [16]). This is consistent with on-off antagonistic receptive fields (see section 4.2). Surround inhibition was found for patch gratings of various spatial frequencies in a large portion of recorded V1 cells, but with high variability between layers [64] and cell types [115]. Indications were also found that suppression may partly arise from network effects within V1, in the form of inhibition from Sst neurons [115]. One study used flashing spots to map receptive fields, finding that 32 % of recorded cells had circular fields with an antagonistic surround [117]. The study further found that, for so-called simple cells (17%), most had fields with elliptical, antagonistic sub-regions [117]. That is, the side regions inhibited the center region's responses to flashing light [117]. The remaining cells showed no clear on and off-regions and responded mainly to moving bars [117]. A summary of the experimental findings of the mentioned studies are shown in tables 7.4 and 7.3.

7.1.3 Effect of anaesthesia

Most studies conducted on the mouse LGN have used anaesthetised mice, while only a few experiments have used awake mice [16,83]. However, the studies which have been presented in this section imply that suppression is at least equally strong in awake mice. In order to determine to which degree the LGN and V1 results in anaesthetised mice is representative of awake mice, Durand et al. [16] did a comparison. They found little effect on spatial processing. Specifically, the preferred spatial frequency and receptive field area was not significantly altered. Nevertheless, it is possible that the spatial tuning curve in the anaesthetic state is scaled by a factor compared to the awake state, in a way that does not influence the optimal frequency. This has also been found in V1 neurons in rabbits [119].

For size tuning measurements, Self et al. [64] found a strong reduction in surround suppression with strength of anaesthesia when measuring neurons in V1. In LIV of V1, reducing the strength of anaesthesia increased the median suppression index from 0.04 to 0.37 [64]. This is consistent with findings by a Vaiceliunaite et al. [82], who compared V1 neurons in awake vs anaesthetised mice. They found that neurons in awake mice displayed the strongest surround suppression [82]. Specifically, 65% of cells in awake mice demonstrated significant surround suppression, against 47% in anaesthetised mice [82]. Additionally, the strength of suppression, classified by the suppression index (SI), was higher for awake mice (SI = 0.70) compared to anaesthetised mice (SI = 0.45) [82]. Adesnik et al. [115] also found that anaesthesia had an effect on size tuning in V1. Neurons in awake mice were significantly more suppressed (SI = 0.93) than the anaesthetised mice (SI = 0.33) in response to large stimuli [115]. One can therefore assume that the findings of surround suppression in anaesthetised mice is not caused by, but rather limited by the anaesthesia, and that the behaviour is also to be found in awake mice.

7.2 Temporal Tuning in the Mouse

7.2.1 LGN

Durand et al. [16]

As well as looking at spatial frequency tuning, Durand et al. [16] mapped the temporal frequency tuning of LGN cells in awake and anaesthetised adult mice. The same neurons were used as for spatial frequency tuning (456 in awake mice, 164 in anaesthetised mice). The visual stimuli were drifting gratings of various temporal frequencies, spatial frequencies, contrasts and orientations, where the optimal combinations of spatial settings were used to determine the preferred temporal frequencies. The mean optimal frequency was 6.27 Hz in awake mice vs 4.07 Hz in anaesthetised mice. The gratings preferred by most cells had frequencies of 8 Hz and 4 Hz in the two states respectively. For all tested frequencies (1, 2, 4, 8, 15 Hz), there were at least some cells responding optimally, indicating that preferred frequencies in LGN span from at least 1 to 15 Hz. [16]

Grubb and Thompson [113]

In addition to their spatial tuning measurements, Grubb and Thompson [113] recorded temporal tuning of LGN cells in adult anaesthetised mice. The temporal tuning was measured using full-field drifting vertical gratings of 70% contrast and the optimal spatial frequency of the cell. The optimal temporal frequency was calculated in 47 cells. The cells showed band-pass tuning, where most cells preferred frequencies around 4 Hz, although preferences ranged from about 1 to 10 Hz. The mean recorded optimal frequency was 3.8 Hz. [113]

Tang et al. [28]

Tang et al. [28] also looked at temporal frequency tuning in addition to their spatial tuning measurements. They recorded from the LGN of adult anaesthetised mice, and the temporal tuning of 58 cells was determined. The stimuli used were sinusoidal gratings of 0.03 cpd spatial frequency and 98% contrast, with temporal frequencies ranging from 0.3 to 9.6 Hz. Most cells (85%) displayed band-pass tuning, and in this group, preferred frequencies ranged from 2.0 to 5.5 Hz, with a median of 3.2 Hz. Out of the remaining cells, neurons with both low-pass and high-pass tuning were found. [28]

Summary of findings

Three studies looking at temporal frequency tuning in the mouse LGN have been presented. Durand et al. [16] looked at both awake and anaesthetised mice, finding mean preferred frequencies of 6.27 Hz and 4.07 Hz respectively. Grubb and Thompson [113] and Tang et al. [28] recorded only from anaesthetised mice, and reported mean peak frequencies of 3.9 Hz [113] and 3.2 [28]. The results in relation to anaesthesia will be discussed in section 7.2.3. Key points from the three studies are listed in table 7.5.

study	stimuli	state	pref tf (Hz)
Durand et al. [16]	full-field grating	awake & anaesthetised	mean: 6.27 (awake) vs 4.07 (anaesthetised)
Grubb & Thompson [113]	full-field grating	anaesthetised	mean: 3.8
Tang et al. [28]	full-field grating	anaesthetised	mean: 3.2 [a]

Table 7.5: Experimental studies on temporal frequency tuning in the mouse LGN.

[a] This mean value was calculated only using cells with band-pass tuning (85% of cells).

7.2.2 V1

Durand et al. [16]

Besides previously mentioned measurements, Durand et al. [16] mapped the temporal frequency tuning in all layers of V1 in awake and anaesthetised adult mice. The same cells as used for spatial tuning were recorded, which were 232 and 80 cells from the two states respectively. The visual stimuli were drifting gratings of various temporal frequencies, spatial frequencies, contrasts and orientations. Optimal combinations of spatial settings were used to determine preferred temporal frequencies. The mean preferred frequencies were 2.99 Hz (awake) and 2.8 Hz (anaesthetised). In both states, most cells responded optimally to the lowest frequency (1 Hz). However, a number of cells showed peak response even for the maximum frequency of 15 Hz. [16]

Van den Bergh et al. [116]

Van den Bergh et al. [116] measured temporal frequency tuning in 69 neurons from all layers of V1 in adult, anaesthetised mice. The optimal temporal frequency was found by using drifting gratings as stimuli. Temporal frequencies were varied from 2 to 15 Hz, with orientation and spatial frequency tuned to each cell's preferred values. The mean and median preferred frequencies were reported as 3.4 and 3.79 Hz respectively. A large proportion of cells with both low-pass and band-pass tuning were found. [116]

Andermann et al. [120]

Andermann et al. [120] mapped temporal frequency tuning in LII/III of V1 in awake adult mice. Patches of drifting gratings were used as visual stimuli. The temporal tuning was obtained from 87 responsive V1 cells. Seven temporal frequencies of 0.5 to 24 Hz and five spatial frequencies from 0.02 to 0.32 cpd were used to determine the optimal values. A broad distribution of preferred temporal frequencies was found, covering the full range of tested values. The median optimal frequency was 3 Hz. [120]

LeDue et al. [121]

LeDue et al. [121] recorded preferences for temporal frequency in anaesthetised adult mice. 188 responsive cells were selected for analysis. The visual stimuli were patch gratings of full contrast, using each cell's optimal orientation and drifting direction. Spatial and temporal frequencies were in the ranges 0.01 - 0.32 cpd and 0.25 - 8 Hz. The mean preferred temporal frequency was 1.77 Hz, with a broad distribution spanning the whole range of presented frequencies. [121]

Gao et al. [122]

Gao et al. [122] measured V1 tuning preferences in 10-14 weeks old anaesthetised mice. Responses were recorded in 220 cells sampled from LII/III to LVI. Drifting sinusoidal gratings were used as stimuli. Temporal tuning measurements were made using maximum contrast gratings of optimal orientation and a spatial frequency of 0.03 cpd. Temporal frequencies were then varied from 0.5 to 16 Hz. 90% of cells showed significant temporal selectivity, out of which 44% showed low-pass tuning and 56% showed band-pass tuning. The low-pass and band-pass populations had significantly different median preferred frequencies of 1.2 and 1.9 Hz respectively. [122]

Niell and Stryker [123]

Niell and Stryker [123] measured tuning properties of V1 in anaesthetised adult mice. Temporal frequency tuning measurements were made using sinusoidally contrast-reversing gratings of optimal orientation, with a spatial frequency of 0.04 cpd. The response to temporal frequencies of 1, 2, 4 and 8 Hz were then recorded. The overall median preferred frequency was 1.68 Hz, and most cells responded optimally to the frequency of 2 Hz. A significantly different temporal tuning behaviour was however displayed by neurons in LIV, where the preferred frequencies had a median of 2 Hz and mean around 4 Hz. [123]

Summary of findings

Six studies have been presented which measured temporal frequency tuning in the mouse V1. A large variation between cells is found, with cells preferring a range of different frequencies. Gao et al. [122] found about equal amounts of low-pass and band-pass tuned cells. Niell and Stryker [123] reported a significant difference in tuning between layers, with LIV preferring higher frequencies. Preferred frequencies were also found to vary between studies, as presented in table 7.6. This could have been influenced by the choice of grating parameters, such as the temporal frequencies presented. For example, the relatively high frequencies reported by Van den Bergh et al. [116] could have been influenced by only presenting frequencies down to 2 Hz, which is higher than some of the mean [121] and median [122, 123] frequencies reported by other studies. Another cause could for example be different sampling from the various layers between the studies.

7.2.3 Effect of anaesthesia

Durand et al. [16] compared the temporal tuning in awake and anaesthetised mice. They found no statistically significant difference in preferred frequencies between the states in V1 [16], consistent with previous findings in rabbit V1 [16, 119]. In LGN however, there was a significant difference between preferred frequencies: 6.27 Hz for awake mice vs 4.07 Hz in anaesthetised mice. This difference is also consistent with

study	stimuli	state	pref tf (Hz)
Durand et al. [16]	full-field gratings	awake & anaesthetised	mean: 2.99 (awake) vs 2.8 (anaesthetised)
Van den Bergh et al. [116]	full-field gratings	anaesthetised	mean: 3.4 median: 3.79
Andermann et al. [120]	patch gratings	awake	mean: 3
Ledue et al. [121]	patch gratings	anaesthetised	mean: 1.77
Gao et al. [122]	full-field gratings	anaesthetised	median: 44% of cells: 1.2 56% of cells: 1.9
Niell & Stryker [123]	full-field gratings	anaesthetised	median: 1.68

Table 7.6: Experimental studies on temporal frequency tuning in the mouse V1.

recordings from rabbit LGN [16, 124]. When later comparing temporal tuning in the Allen LGN to experiments, most emphasis will therefore be put on the findings by Durand et al. [16] in the awake mouse. For V1 on the other hand, the state will not be considered when referring to experiments.

Chapter 8

Simulation Method

8.1 The Versions of the Allen Model

Simulation measurements were run using different versions of the Allen LGN model. The first is the original model, where the spatial filters of the LGN ON and OFF cells are simple 2D Gaussians. The second will be referred to as the M1 LGN, where the Gaussians have been replaced with DoG filters, given by equation (5.6), with parameters $A_c = A_s = 1.73$ and $\sigma_s = 2.45\sigma_c$. In addition, four DoG filters with other parameters are used, which will be collectively denoted by M1$*$. The parameters for these filters are given in section 8.1.2. A qualitatively different model, denoted M2, refers to an LGN where firing rates have been normalised, but spatial filters remain unchanged. In all the above cases, it is only the LGN which has been changed, not V1. However, to easier separate between them, the corresponding V1 will be referred to as the original, M1, M1$*$ or M2 models, depending on which LGN is used to drive the network. For certain measurements, the M1 V1 model was run without recurrent connections between V1 cells, which will be called the M1_{norec} model. An overview of the different modifications are given in table 8.1.

Model	Description
Original	The original model.
M1	LGN spatial filters are DoGs from equation (5.6), where $A_c = A_s = 1.73$ and $\sigma_s = 2.45\sigma_c$.
M1$*$	Similar to M1, but with four new sets of DoG parameters (see list in section 8.1.2).
$\mathbf{M1_{norec}}$	As in M1, but without recurrent connections in V1.
M2	LGN responses are normalised via equations 8.1 and 8.2.

Table 8.1: The versions of the Allen model used to run simulations. For the M1, M$*$ and M2 models, changes were made only to the LGN, while for M1_{norec} both LGN and V1 has been modified.

8.1.1 The M1 model

The spatial filters of the Allen LGN were modified so that each single Gaussian was replaced by a Difference-of-Gaussians (DoG), as given by equation (5.6). The modified script defining the DoG filters is available via GitHub[1].

Choice of DoG parameters

A neuron in the original LGN model has a receptive field size defined by σ, drawn from a triangular distribution with an experimentally supported range [12,16]. The spatial filter was modified by adding an inhibitory surround, keeping the single Gaussian as the center component. Specifically, the equation for the spatial filter was changed from 5.4 to 5.6, where the parameters needing to be specified are

- A_c: amplitude of the center Gaussian
- A_s: amplitude of the surround Gaussian
- σ_c: spread of the center Gaussian
- σ_s: spread of the surround Gaussian

The amplitude of the center and surround Gaussian, A_c and A_s, were set equal in order to obtain maximum surround suppression, since the aim was to determine whether surround suppression could be observed in the model. The two amplitudes were scaled by the same constant, so that the firing rate for the preferred spatial frequency of the DoG model matches the firing rate for the preferred spatial frequency of the original LGN model. The scaling factor was found by comparing the analytical spatial frequency tuning curves of the original and DoG spatial filter, as illustrated in figure 5.2.

The parameter σ of the single Gaussian became σ_c of the center in the DoG filter. There was then only one free parameter left, the sigma of the surround Gaussian function, σ_s, which was set to $C\sigma_c$, where C is a fitting constant. The distribution of preferred spatial frequencies found by Durand et al. [16] was used as a reference when determining the value of C. Specifically, C can be set so that the median preferred frequency matches the experimentally measured median of 0.08 cpd [16]. This was achieved by calculating the optimal spatial frequencies of all the LGN cells, as described below.

Calculating optimal spatial frequencies of LGN cells

To find the preferred spatial frequencies of LGN cells, the spatial frequency tuning of each cell was calculated using equation (5.1). The spread of the surround Gaussians was determined by setting $\sigma_s = C\sigma_c$, where C was a variable to be adjusted. The values of σ_c ($= \sigma$) for each cell were obtained from the parameters listed in the file `lgn_full_col_cells_3.csv`[2] in the LGN model. This file gives a parameter named the "spatial size" of each cell , equal to 3σ [111]. Hence, σ_c was set equal to "spatial size"/3. With the values of σ_s and σ_c determined, equation (5.1) was used to calculate the spatial frequency tuning of each cell. The optimal frequency was taken as the peak of the tuning curve. The median optimal frequency of all LGN cells

[1] The code for creating the DoG filters is available at `https://github.com/lenamyk/Exploring_the_Allen_model/blob/master/adapted_from_Allen_Institute/LGN_modifications`

[2] See [112] for link to `lgn_full_col_cells_3.csv`.

was then calculated. The constant scaling factor, C, was subsequently adjusted to obtain a median optimal spatial frequency of 0.08 cpd, the same as found by Durand et al. [16]. The Python script used for calculating peak frequencies and finding the median is given in the appendix C.1.

8.1.2 The M1$*$ models

In addition to the M1 model, four different DoG filters were tested. The new model variations were used for analysing size tuning in V1 in response to patch grating stimuli. Stimuli parameters are given in table 8.2c. The aim was to check whether the lack of suppression in the M1 model was consistent across the models, as opposed to the result of the chosen DoG parameters. The parameters of the new filters are

- $A_c = A_s = 1.2$ and $\sigma_s = 5\sigma_c$
- $A_c = A_s = 1.06$ and $\sigma_s = 10\sigma_c$
- $A_c = A_s = 1$ and $\sigma_s = 50\sigma_c$
- $A_c = 1.5$, $A_s = 0.5A_c$ and $\sigma_s = 2.45\sigma_c$.

As before, the amplitudes have been scaled so that the peak activity corresponds to the peak activity in the original LGN model. The scaling factors were determined by comparing analytical spatial frequency tuning curves of the selected filter with the original model (see example illustration in figure 5.2).

8.1.3 The $\mathbf{M1_{norec}}$ model

The response of the M1 V1 model was compared to the same model, but without recurrent connections: $\mathrm{M1_{norec}}$. The motivation for this was to look into the effect of recurrent connections on size tuning, since these connections have been suggested to be involved in surround suppression [115,125]. The recurrent connections were turned off by removing the files[3] specifying the connectivity between V1 cells from the list of V1 inputs in the file `config.json`[4] .

8.1.4 The M2 model

To model nonlinear suppression in the LGN, the activity of each cell per time step was normalised based on the net LGN activity per time. The modified LGN script which includes this normalisation is uploaded on GitHub[5]. The M2 LGN was then driven with patch gratings (see table 8.2 of stimuli parameters) to measure its size tuning. The output was further used to drive the V1 model, to see how nonlinear suppression in the LGN may influence the size tuning of V1 cells.

Starting with the original Allen model, the LGN was modified so that the normalised firing rates for each cell are given by

$$r_{\mathrm{norm},\,i}(t) = \alpha \frac{r_i(t)}{\beta + (\frac{1}{N}\sum_{k=0}^{N} r_k(t))} \tag{8.1}$$

[3]V1 recurrent connection are specified in `v1_v1_edges.h5` [126] and `v1_v1_edge_types.csv` [127].

[4]See [128] for link to `config.json`.

[5]The modified LGN script which normalises the final firing rates, `lgn_functions_normalise_rates.py`, is available at: `https://github.com/lenamyk/Exploring_the_Allen_model/blob/master/adapted_from_Allen_Institute/LGN_modifications`

where $r_i(t)$ is the unnormalised firing rate of cell i at time t and N is the number of LGN cells to sum over. The form of the equation is inspired by [129]. The parameter β was set equal to one. The number of cells, N, here includes all cells, but may be changed to include only specific cells, based on for example proximity or type. The scaling parameter α was chosen so that the spontaneous rates of the cells are unchanged, that is

$$\alpha = \beta + (\frac{1}{N}\sum_{k=0}^{N} r_{\text{spont},\,k}) \tag{8.2}$$

$$= 4.84 \tag{8.3}$$

An illustrative example of an LGN cell before (black) and after (blue) being normalised is shown in figure 8.1, showing that both models have the same spontaneous activity (grey) and a similar peak rate.

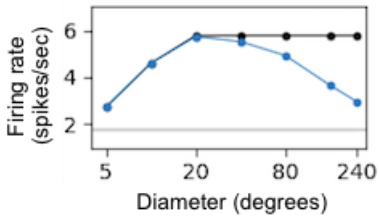

Figure 8.1: Size tuning of an sON-TF8 neuron in the M2 LGN (blue) and the original LGN (black) in response to patch gratings with temporal frequency of 4 Hz and spatial frequency of 0.04 cpd. The spontaneous rate of both models (overlapping) is marked by the grey line.

8.2 Running the Models

Since the Allen model is still under development, a few modifications had to be made to the LGN module before running it. A slightly modified version of `nwb.py` from the Brain Modeling Toolkit (BMTK) [130] was substituted and imported instead of `see_engine.nwb`, since the latter, which was needed for spike train generation, was not found amongst the LGN scripts. The modified `nwb.py` file is uploaded on GitHub as `nwb_copy.py`[6]. In addition, the LGN scripts were updated to be compatible with Python 3. The model is run via `simulate_drifting_gratings.py`[7], which calls the functions for creating visual input, calculating firing rates and generating spike output.

The spike train output from the LGN module were used as input for the V1 model. The main simulation parameters and input for running the V1 model are specified in the configuration file `GLIF_network/config.json`[8]. This allows turning on and off input from the LGN, background noise and recurrent connections. The V1 simulations were then run from the script `run_pointnet.py`[9].

[6] A modified copy of the `nwb_copy.py` is available at: `https://github.com/lenamyk/Exploring_the_Allen_model/tree/master/adapted_from_Allen_Institute/LGN_modifications`

[7] See [131] for link to `simulate_drifting_gratings.py`.

[8] See [128] for link to `config.json`.

[9] See [132] for link to `run_pointnet.py`.

8.3 Visual Stimuli

Incorporated in the Allen LGN model is a script that creates full-field drifting gratings, `movie.py`[10]. The visual stimuli created by modifying this script are available through GitHub[11]. The following sections describes the various visual stimuli used to drive the Allen model, how they were created and the parameters that where chosen. For simulations where spatial filter responses were analysed, the parameters are listed in table 8.2. For temporal frequency tuning, the parameters are listed in table 8.3. Common for all simulations is that they have a 3-second duration and a time step of 0.001 seconds (1 ms).

Creating drifting gratings

The grating stimuli used were sinusoidally varying drifting gratings. Before running different grating simulations, various parameters were specified, including

- Orientation of the grating
- Total duration of visual input
- Duration of initial grey screen
- Spatial resolution
- Spatial frequency
- Temporal frequency
- Intensity contrast

For full-field drifting gratings and patch gratings, a gray screen input is presented during the first 0.2 seconds. The gray screen frames were initially used in order to let the network stabilise, however this was not found to make much difference. Still, the gray screen was kept to make it easier to identify spontaneous activity and calculate the firing rates of LGN cells. The contrast for both types of gratings was set to 1 (100%).

Full-field gratings

Full-field gratings, like the one shown in figure 8.2, were used for spatial and temporal frequency tuning measurements. Horizontally drifting gratings with vertical bars (orientation 0°) were used, corresponding to equation (5.7). These simulations were originally only for analysing LGN cells, where the majority of cells are either ON or OFF with no orientation preference. For the single cell analyses, only these cells were examined. Hence the orientation was arbitrarily chosen.

When measuring spatial frequency tuning, the temporal frequency was set to 4 Hz, which is the median preferred temporal frequency of the LGN and V1 neurons (see section 10). Measurements were taken for spatial frequencies 0.005, 0.04, 0.08, 0.12, 0.16, 0.20, 0.24, 0.28 and 0.32 cpd. Since frequencies up to 0.32 cpd were used, the gratings used for spatial tuning measurements required higher resolution than the

[10] See [133] for link to `movie.py`

[11] The modified scripts for creating patch gratings and flashing spots are available at: `https://github.com/lenamyk/Exploring_the_Allen_model/tree/master/adapted_from_Allen_Institute/LGN_modifications`

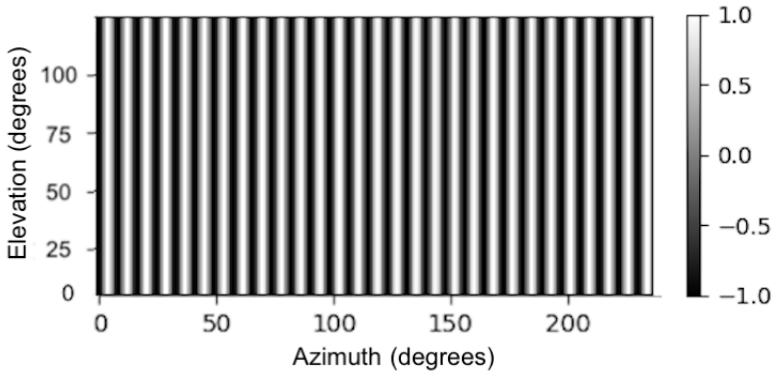

Figure 8.2: A full-field grating of spatial frequency 0.12 cpd, as used to drive the LGN model. The dimensions of the input are $240° \times 120°$. Visual coordinates are measured from bottom left. Azimuth and elevation refers to the horizontal and vertical angular coordinates respectively.

other stimuli. It was found that a resolution of 2 pixels per degree was sufficient to give a fairly good fit to the analytical spatial frequency tuning curve, as will be seen in section 9.2.4. This resolution was therefore chosen for the final measurements.

For temporal frequency tuning in the original model, the spatial frequency was set to 0.04 cpd. For temporal frequency tuning in M1, the spatial frequency was set to 0.08 cpd, which is the median preferred spatial frequency of the M1 LGN (section 9.2.3) and M1 V1 (section 9.2.4) which were used to run the measurements. The temporal frequencies used were 1, 2, 4, 8, 15 and 30 Hz. All frequencies were selected so that responses over an integer number of oscillation periods could be measured. These frequencies also include those used to record temporal tuning by Durand et al. [16], the study which has been used to optimise the Allen LGN cells [12].

Patch gratings

Circular patches of moving gratings were used for measuring size tuning in the original and modified Allen model. The stimuli were created using the same input script as for full-field drifting gratings. In order to create patches of different diameters, a circular mask of a specified radius was applied to the grating, using the standard circle equation

$$r^2 = (x - x_0)^2 + (y - y_0)^2 \tag{8.4}$$

where $(x_0 = 120, y_0 = 60)$ is the spacial center of the input matrix. Diameters of 5°, 10°, 20°, 40°, 80°, 160° and 240° were used, where 240° is the width of the visual input frame. Note that the patches with diameters 160° and 240° are therefore not fully circular, but restricted by the maximum height of the input.

The orientation of the patch gratings was set to 155.26°. This angle was chosen as it is the preferred orientation of the most centered V1 neuron, which was the first neuron to be analysed. The aim was to ensure that the stimuli evoked a strong response in the analysed neuron. Multiple other neurons near the center were also found to respond strongly to this stimulus, even though the orientation was not their preferred. Hence, the measurements were not repeated for any other orientations, and the angle was kept at 155.26 ° for all patch grating simulations.

The resolution for the patch grating stimulus was set automatically, depending on the spatial frequency, by the input script provided by Allen LGN model. This is the

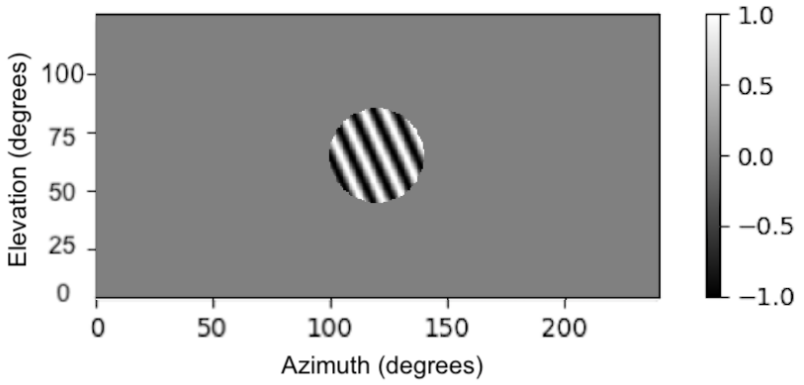

Figure 8.3: One of the patches of drifting grating used to drive the LGN model. The spatial frequency is 0.12 cpd, and the grating is aligned along the preferred orientation of the most centered V1 neuron, at 155.26°. The dimensions of the input are 240°×120°, and the patch has a diameter of 40°. Visual coordinates are measured from bottom left. Azimuth and elevation refers to the horizontal and vertical angular coordinates respectively.

resolution needed to visualise all the gratings with the same amount of gray scales. Stimuli with low resolution are also significantly faster to simulate. The temporal frequency for all patch gratings was 4 Hz, which is the median preferred value of the LGN and V1 (chapter 10).

The original, M1∗ and M2 Allen models were driven with patch gratings of 0.04 cpd spatial frequency. For the M1 Allen model, patch gratings of several spatial frequencies were used, in order to evaluate the consistency of the size tuning across frequencies. Initially, measurements were repeated for 0.04, 0.08, 0.12, 0.16 and 0.20 cpd, which covers the range of preferred spatial frequencies found in section 9.2.3. Later, the measurements were also run for 0.02 cpd, since calculations by [95] suggested that more surround suppression is observed for lower spatial frequencies. Comparisons of size tuning between the M1 and M1_{norec} model were made for patch gratings with spatial frequency of 0.08 cpd, which was the preferred frequency of the V1 M1.

Flashing spots

Flashing spots were also used for size tuning measurements in the original and M1 Allen model. The stimulus was composed of an alternation between equal durations of white circles and full-field grey frames. For both the original and modified Allen model, a sequence of 20 spots (75 ms flash duration) was presented during each trial. Additionally, to see if the results varied with flash duration, measurements were repeated with 5 flashes (300 ms duration) per trial in the modified Allen model. First, a sequence of gray and white frames were created. Then, the spots were defined by applying the circular mask given by equation (8.4) to the white full field frames. Like for patch gratings, diameters of 5°, 10°, 20°, 40°, 80°, 160° and 240° were used, where 240° is the full width of the visual input. All spots were centered at the middle of the visual input. The resolution was set to 1 pixel per degree, so that flashing spots down to 5° in diameter could be reasonably accurately visualised in the model.

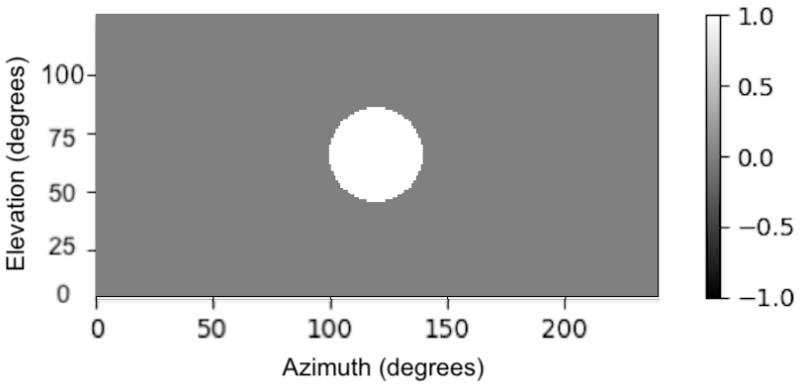

Figure 8.4: A flashing spot of 40° diameter which was used as input to the LGN model. The dimensions of the input are $240° \times 120°$. Visual coordinates are measured from bottom left. Azimuth and elevation refers to the horizontal and vertical angular coordinates respectively.

stimulus	ori (°)	total dur (s)	grey dur (s)	res (pix/°)	sf (cpd)	tf (Hz)	contrast (%)
full-field grating	0	3	0.2	2	0.005 [a] -0.32	4	100
patch grating	155.26	3	0.2	0.4	0.04	4	100
flashing spots	NaN	3	0	1	NaN	NaN	NaN

(a) Original model

stimulus	ori (°)	total dur (s)	grey dur (s)	res (pix/°)	sf (cpd)	tf (Hz)	contrast (%)
full-field grating	0	3	0.2	2	0.005 [b] - 0.32	4	100
patch grating	155.26	3	0.2	0.2 [c] - 2	0.02 [d] - 0.2	4	100
flashing spots	NaN	3	0	1	NaN	NaN	NaN

(b) M1 model

stimulus	ori (°)	total dur (s)	grey dur (s)	res (pix/°)	sf (cpd)	tf (Hz)	contrast (%)
patch grating	155.26	3	0.2	0.8	0.08	4	100

(c) $M1_{norec}$ model

stimulus	ori (°)	total dur (s)	grey dur (s)	res (pix/°)	sf (cpd)	tf (Hz)	contrast (%)
patch grating	155.26	3	0.2	0.4	0.04	4	100

(d) M1∗ models

stimulus	ori (°)	total dur (s)	grey dur (s)	res (pix/°)	sf (cpd)	tf (Hz)	contrast (%)
patch grating	155.26	3	0.2	0.4	0.04	4	100

(e) M2 model

Table 8.2: The parameters used in the simulations for studying spatial properties of the LGN and V1 model, grouped by the type of model. The columns show, from left to right, stimulus, orientation of grating, total duration of visual input, duration of grey screen before stimulus, spatial resolution (in pixels per degree), spatial frequency, temporal frequency and intensity contrast.

[a] Measurements at 0.005, 0.04, 0.08, 0.12, 0.16, 0.20, 0.24, 0.28 and 0.32 cpd

[b] Measurements at 0.005, 0.04, 0.08, 0.12, 0.16, 0.20, 0.24, 0.28 and 0.32 cpd

[c] Measurements at 0.02, 0.04, 0.08, 0.12, 0.16 and 0.20 pix/°

[d] Measurements at 0.02, 0.04, 0.08, 0.12, 0.16 and 0.20 cpd

stimulus	ori (°)	total dur (s)	grey dur (s)	res (pix/°)	sf (cpd)	tf (Hz)	contrast (%)
full-field grating	0	3	0.2	0.4	0.04	1-30 [a]	100

(a) Original model

stimulus	ori (°)	total dur (s)	grey dur (s)	res (pix/°)	sf (cpd)	tf (Hz)	contrast (%)
full-field grating	0	3	0.2	0.8	0.08	1-30 [b]	100

(b) M1 model

Table 8.3: Parameters of the visual stimuli used for temporal frequency tuning in the original and M1 model. The columns show, from left to right, stimulus, orientation of grating, total duration of visual input, duration of grey screen before stimulus, spatial resolution (in pixels per degree), spatial frequency, temporal frequency and intensity contrast.

[a] Measurements at 1, 2, 4, 8, 15 and 30 Hz

[b] Measurements at 1, 2, 4, 8, 15 and 30 Hz

8.4 Calculating Firing Rates

8.4.1 Time average vs average modulation

Depending on the linearity of a neuron, drifting grating stimuli may evoke different types of responses in the cell. The two responses are quantified by the time averaged mean firing rate (F0) and the mean modulation amplitude (F1) [134]. Both are commonly used measures for quantifying neuron responses in model cells and in the real mouse [12, 16, 52, 134].

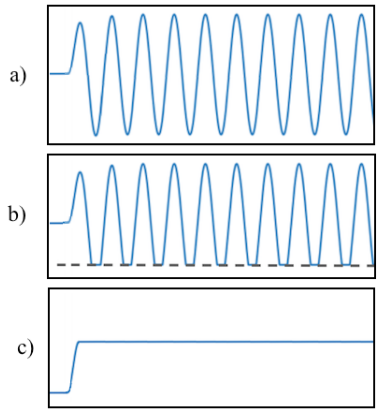

Figure 8.5: Firing rates for three types of neurons in response to a grating stimulus. (a) A hypothetical, purely linear neuron will respond to grating stimuli by a firing rate modulation with the same temporal frequency as the grating [52]. (b) An LNP neuron such as an Allen LGN cell responds to grating stimuli by a firing rate modulation, where negative firing rates are set to zero by a threshold (dashed line) [12]. (c) A nonlinear neuron responds to grating stimuli by an increase in mean firing rate, but with little or no modulation [135].

For a hypothetical, completely linear response, the activity will oscillate around a constant spontaneous rate, with modulations having the same temporal frequency as the grating [52] (figure 8.5a). In this case, the response will only be in terms of F1, with no change in F0 from the base value. A nonlinear neuron, on the other hand, will react to the drifting grating stimuli by an increase in F0, while the firing rate as a function of time is mainly constant during the grating stimulation (figure 8.5c) [135]. For neurons such as the Allen LGN cells, which have linear responses rectified at zero (equation (3.10)), the F1 and F0 tuning responses usually follow each other [16, 134]. This is because the rectification, which prevents negative rates, causes the mean firing rate to increase when the rate's modulation amplitude increases.

Figure 8.6 demonstrates this result. The plot shows the spatial frequency tuning curves of an M1 LGN neuron, measured in terms of the F0 and F1 responses. For low activities, the modulation amplitudes are too small to reach the zero-threshold of the rectifier. The F0 response is therefore just the spontaneous rate, while the amplitude is positive (see rate at 0.16 cpd in figure 8.6). For higher activity, the two tuning curves have similar shapes, with the same peak frequency.

In the main result sections of this thesis, activity in response to grating stimuli will, as with flashing spots, be referred to in terms of the total mean firing rate, F0. The exception is when comparing LGN spatial frequency tuning to the analytical solutions in equations (5.8) and (5.9). These equations give the spatial frequency

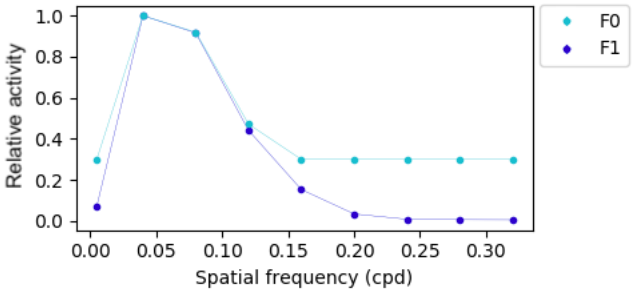

Figure 8.6: Spatial frequency tuning of an M1 LGN neuron, quantified by the normalised F0 (cyan) and F1 (dark blue) responses. The stimuli is a full-field drifting grating with temporal frequency of 4 Hz.

dependence of the F1 response, and will therefore be compared with the simulated F1 response. Comparisons between the F0 and F1 responses for one example LGN and V1 neuron from each population are shown in the appendix, for size tuning with patch gratings (B.8, B.9, B.26, B.27), as well as spatial (B.4, B.5, B.14, B.15) and temporal (B.44 B.46) frequency tuning with full-field gratings. Tuning curves include both the original and M1 model. In all cases, the F1 and F0 tuning responses were found to follow each other.

8.4.2 LGN firing rates

In the LGN model, the visual input is transformed and firing rates are calculated for each neuron [12]. Different trials of spike trains are then generated by a Poisson process based on the firing rates [12]. For all measurements of the LGN response, the raw firing rate output was used directly, instead of going via spike trains and converting this back to a firing rate. Since the LGN firing rates are generated deterministically in the model, this avoids introducing the measurement uncertainty caused by the stochastic spike generation. The output from the LGN model includes firing rates for each neuron for the simulated time-duration. The output for an example neuron is shown in figure 8.7 and 8.8 for drifting grating and flashing spot stimuli. For an example of the code used to analyse LGN firing rates, see section C.1 in the appendix.

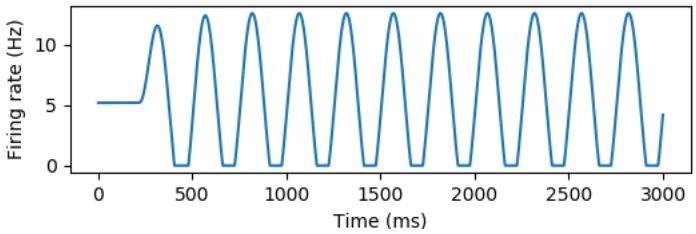

Figure 8.7: Example firing rate for an example LGN neuron (id = 8018). The simulation duration is 3 seconds, with 0.2 seconds initial grey screen. The visual stimuli is a full-field drifting grating with spatial frequency of 0.08 cpd and temporal frequency of 4 Hz.

The time averaged response was taken as the mean firing rate over an interval of periodic activity. The rate was measured starting at half the simulation time (1500

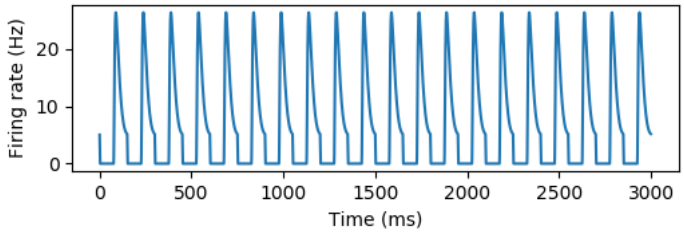

Figure 8.8: Example firing rate for an example LGN neuron (id = 8018). The stimuli is a sequence of flashing spots of 20 degrees diameter and 75 ms durations. The total length of the simulation is 3 seconds.

ms), to ensure that oscillations had stabilised. For grating stimuli, firing rates were measured over a 1000 ms interval. This corresponds to an integer number of periods for all temporal frequencies used: 1, 2, 4, 8, 15, and 30 Hz. For flashing spots, the time averaged activity was measured over a 1200 ms interval, which corresponds to the onset of one flash to one later in the simulation.

The F1 response to grating stimuli was found by taking the discrete Fourier transform (DFT) of the firing rate. The DFT can be defined as [136]

$$X(k) = \sum_{m=0}^{N-1} A(m) e^{2\pi i j m/N}; \qquad k = 0, 1, ..., N-1 \tag{8.5}$$

where $A(m)$ are complex-valued Fourier coefficients and $N = T/dt$ the number of samples. The amplitudes as a function of frequency are then given by the normalised absolute values of the transform, $2N^{-1}\,|X(k)|$, for $k = fNdt$ [137]. The time step dt, used for all simulations was set to 0.001 seconds. The measurement window, of length T, was the same as when computing the time average, 1000 ms. The DTF of the firing rate was computed with the Fast Fourier Transform (FFT) algorithm [138], using the `numpy.fft.fft` function in Python. The amplitude spectrum of an example LGN neuron as a function of frequency is shown in figure 8.9. The frequencies corresponding to each element of the amplitude spectrum were extracted via the `np.fft.fftfreq` function in Python. The amplitude spectrum was then evaluated at

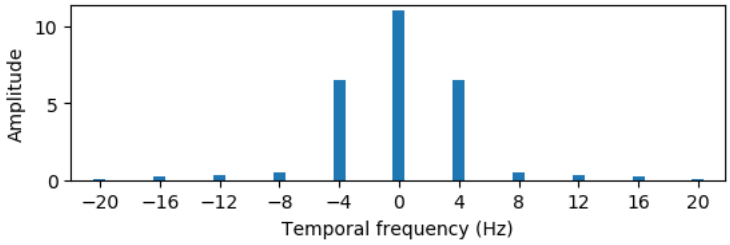

Figure 8.9: Amplitude spectrum for an example LGN neuron in response to full-field gratings with temporal and spatial frequencies of 4 Hz and 0.08 cpd respectively. The F0 response is the zero-frequency in the middle of the plot, while the F1 response is the amplitude at the frequency of the grating, at 4 Hz.

the frequency of the grating stimuli to get the F1 response. By evaluating the zero-frequency component of the amplitude spectrum, one can also find the time averaged rate, F0. This was confirmed to give the same value as when measuring the mean rate over an oscillation period.

8.4.3 V1 firing rates

For V1 neurons, firing rates were averaged over 20 trials of 3-second simulations. For each V1 trial, a different LGN trial was used as input. Hence for each average calculated, there were 20 trials of spike generation from the LGN and a corresponding 20 trials of V1 simulation. Section C.1 shows an example of the code which was used to calculate and visualise firing rates in V1 cells.

The time averaged activity for a V1 neuron was found from the firing rate as a function of time, by calculating the mean rate over the stimulation duration. To obtain firing rates as a function of time, the spike times for the specific neuron were first extracted from the V1 output. The spike train was then converted to a discrete firing rate, by counting the number of spikes per time step, and multiplying by the number of time steps per seconds (1000). An example of the discrete firing rate as a

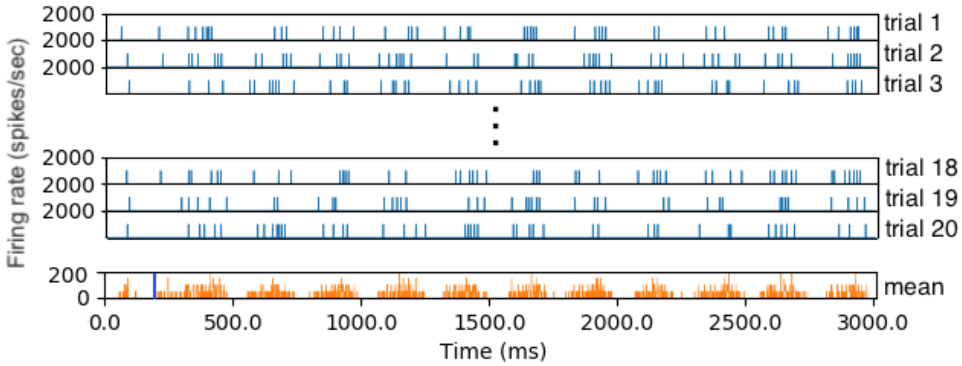

Figure 8.10: Spike trains for an example V1 neuron. The mean firing rate is calculated over 20 trials. Blue curves show discrete firing rates per trial, and the orange curve show the trial average. The visual input used is a full-field drifting grating of 4 Hz temporal frequency. The blue vertical line in the bottom row shows the grating stimuli onset.

function of time for drifting grating stimuli is illustrated in figure 8.10. When looking at a single trial, the spike times seem to be almost random. However, the figure shows that the mean rate over 20 trials oscillates periodically as a function of time, following the temporal frequency of the grating (4 Hz). The time average for each trial was found, and the total mean over all trials was then calculated.

The mean F1 response was taken as the average modulation over all trials. The modulation amplitude was found for each trial by computing the FFT of the rate, again by using the `numpy.fft.fft` function in Python, and evaluating the amplitude spectrum at the frequency of the grating. For V1 rates, the time window used to calculate the F1 response was set from 500 ms to the end of simulation (3000 ms). The initial 500 ms were excluded to make sure that the oscillation amplitude had stabilised. The amplitude spectrum from a single trial for an example V1 neuron is shown in figure 8.11. The amplitude for each trial was found, and the F1 response was taken as the mean amplitude over trials.

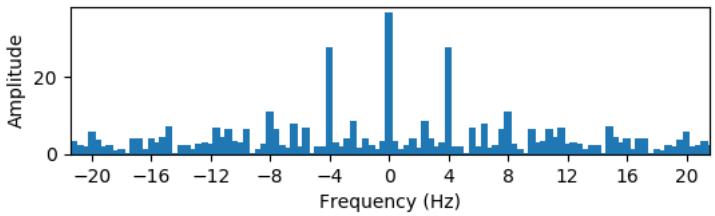

Figure 8.11: Amplitude spectrum for an example V1 neuron in response to full-field gratings with temporal and spatial frequencies of 4 Hz and 0.04 cpd respectively. The zero-frequency component shows the time averaged neuron response. The mean modulation amplitude of the firing rate was was evaluated at the frequency of the grating, at 4 Hz.

8.5 Spatial Frequency Tuning

LGN

To create spatial frequency tuning curves from LGN simulations, full-field drifting gratings of increasing spatial frequencies (see table 8.2) were used as visual input. The measurements were first run for the original model. After having determined the parameters of the DoG filter for the M1 model, the simulations were repeated in the M1 LGN. Tuning curves of the F1 responses were compared with the analytical F1 tuning results from equations 5.8 and 5.9. Neurons with minimum, maximum and midsized receptive fields were considered. This way, tuning curves covering the whole range of preferred frequencies of the LGN cells could be compared to the analytical results. Since equations 5.8 and 5.9 only give the shape of the tuning curves, but with no scaling factor, the analytical curve was scaled to best fit the simulation results.

There are two reasons for establishing that LGN output follows the analytical prediction. First, the analytical solution serves as a benchmark for both the original and M1 Allen model. Second, if the spatial frequency tuning of the LGN cells can be found analytically, the effect of modifying the DoG parameters can be determined without having to run numerous simulations. This is used in section 8.1.1 to specify the value of C in $\sigma_s = C\sigma_c$ from equation (5.6). By tuning C, the same median preferred frequency as found by Durand et al. [16] was obtained for LGN cells.

V1

Spatial frequency tuning measurements for V1 neurons were obtained through simulations, using the original and M1 models. Cells were divided into groups based on the 17 V1 classes, and 10 responsive neurons were selected from each group. The mean firing rate at each frequency was then calculated and plotted as tuning curves for each cell. To get an overview of the preferred spatial frequencies of V1 neurons, a distribution of optimal frequencies of a 1000 randomly sampled neurons was also plotted. Additionally, optimal frequencies within each cell class were found by sampling and analysing 50 neurons of each type. The optimal frequencies were found by determining which frequency evoked maximum activity in each cell. Cells which did not spike in response to gratings of any frequency were discarded from the analysis.

8.6 Temporal Frequency Tuning

LGN

The original and M1 LGN model was driven with full-field drifting gratings of different temporal frequencies (table 8.3). One example neuron was first selected from each cell type. The mean firing rate was calculated for each temporal frequency in both the original and M1 model. The results where compared for the two models, confirming that they have the same temporal tuning.[12] The remaining analysis was done for the M1 model. To obtain a distribution of preferred frequencies in the LGN, the frequencies producing peak mean firing rates were found for all LGN neurons. The results were plotted in a histogram. Furthermore, 50 neurons of each class were randomly sampled. The preferred frequencies of the sampled cells were found and presented in histograms according to cell type. The neurons which did not respond to any of the frequencies were omitted from the plots.

V1

Temporal frequency tuning measurements for V1 neurons were obtained through simulations in the original and M1 model. As for the LGN, tuning curves for the original and M1 model were compared, ensuring that the two models have the same temporal tuning. Further analysis was then done for the M1 model. The method for analysing temporal tuning in V1 was the same as that used for spatial tuning (see subsection in 8.5). In summary, tuning curves of the mean firing rates were obtained for 10 neurons of each type. Additionally, the peak temporal frequencies of 1000 V1 neurons were plotted, as well as the peak frequencies of 50 neurons of each class. As before, the non-responsive neurons were not included in the analysis.

8.7 Size Tuning

LGN

To examine if the Allen model displayed size suppression, flashing spot and patch grating stimuli of increasing diameters were used as visual input to the LGN (see table 8.2). LGN firing rates were analysed for the original, M1 and M2 models. Ideally, the stimulus origin should be at the visual center for the examined neuron. Thus, only response of neurons with receptive field centers near the stimulus center were analysed. As described in section 6.2, the LGN neurons are not represented with physical coordinates. Instead, each LGN cell has a location in visual space representing its receptive field center [12]. The coordinates for each LGN cell are defined in the file `lgn_full_col_cells_3.csv`[13], where $x = 120°, y = 60°$ is the center of the visual input. From each class, the neuron with the receptive field center closest to stimulus center was selected for analysis. For white flashing spot stimuli, responses of ON and OFF neurons were compared, seeing as they have different preferences for luminosity. For this, the eight ON and eight OFF neurons with fields closest to the center were selected. To illustrate how the LGN receptive fields are positioned in visual space, the fields of selected ON-neurons are shown in figure 8.12.

[12] Since the LGN filters are spatio-temporally separable, the original and M1 model have the same temporal tuning, as it is only the spatial filter which is different between the two.

[13] See [112] for link to `lgn_full_col_cells_3.csv`.

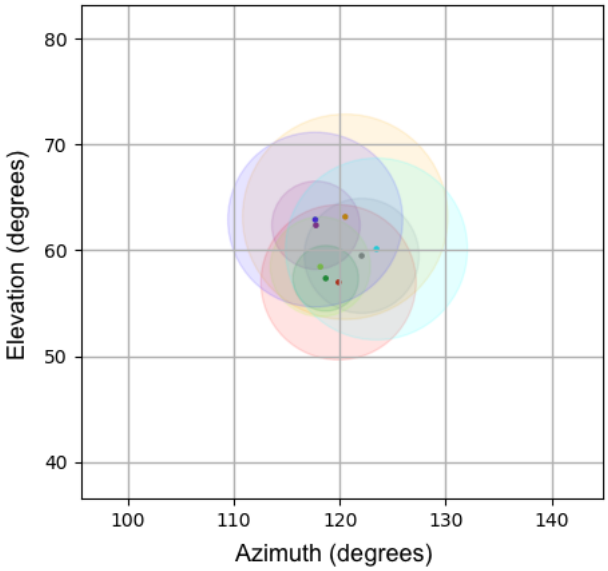

Figure 8.12: Receptive fields of the eight ON-neurons selected for analysing stimulus diameter vs firing rate. The visual center for each neuron is illustrated by the dark dots. The coloured circles represent the full extent of the simple 2D Gaussian receptive fields of the Allen LGN model. Only the middle section of the 240°×120° visual frame is shown. Azimuth and elevation refers to the horizontal and vertical angular coordinates in visual space.

V1

The output activities from the size tuning simulations in the original, M1, M1∗, $M1_{\text{norec}}$ and M2 LGN were used to drive the V1 model, for parameters listed in table 8.2. As in the LGN, the neurons selected for size tuning analysis should have visual centers as close as possible to the stimulus center. Section 6.1 describes the mapping between the receptive fields and somatic coordinates of V1 cells. These somatic coordinates are loaded in from the file `GLIF_network/network/v1_nodes.h5`[14] when running the model. Neurons were selected based on their somatic coordinates as listed in this file. All neurons within a 70 µm radius from the center of cortex (x = 0, z = 0)[15] were located[16]. This corresponds to all neurons with receptive field centers within 4.9° azimuth or 2.8° elevation from the stimuli center. For the original, M1, $M1_{\text{norec}}$ and M2 model, up to ten neurons from each population were selected for analysis, depending on how many were within the 70 µm radius. Size tuning curves where then plotted for all the recorded neurons. For comparisons between the M1 and $M1_{\text{norec}}$ model, six neurons were selected from each population receiving LGN input. These are i1Htr3a, i23Pvalb, i4Pvalb, i5Pvalb, i6Pvalb, e23, e4, e5 and e6. Tuning curves from the two models where then compared for each neuron. For the M∗ models, size tuning of five example neurons within a 70 µm radius were plotted.

[14] See [139] for link to `GLIF_network/network/v1_nodes.h5`.

[15] Somatic coordinates of V1 are given by x, y, z, where x, z spans the cortical plane, while y measures cortical depth.

[16] Script for selecting neurons for size tuning analysis: `https://github.com/lenamyk/Exploring_the_Allen_model/blob/master/code_and_analysis/select_neurons_for_size_tuning.py`

Chapter 9

Size and Spatial Frequency Tuning of the Allen Model

This chapter looks at the spatial response properties of the Allen model, in particular size tuning and spatial frequency tuning. The model is driven by stimuli similar to those used in the experiments referenced in chapter 7: white flashing spots, circular patches of drifting grating and full-field drifting gratings. Initially, the properties of the original Allen model are explored. LGN and V1 simulation results are then compared to the experimental findings of chapter 7.

Chapter 7 referred to several studies on mice LGN and V1. The studies reported that band-pass spatial frequency tuning was observed in both LGN [16, 28, 52, 83, 113] and V1 [16, 116, 118]. In LGN, surround inhibition in response to isoluminous flashing spots was found [52, 83, 113]. Size tuning with patch gratings [48, 64, 82, 115, 116] and flashing spots [117] also showed surround suppression in V1. No experimental studies with patch gratings stimuli were found for mice LGN. Hence, for patch gratings, only the V1 response is compared directly to empirical results.

Simulations from the original Allen model show no size suppression and mainly low-pass spatial frequency tuning in LGN and V1. These results are in contrast to the experimental findings mentioned above. The next step is therefore to explore modifications for making the model more realistic in terms of these properties. As discrepancies are found for both LGN and V1, the approach is to first look at the LGN, to see how changes might affect both the LGN and V1.

One feature separating the Allen LGN from that of real mice is the center-only organisation of the cells' receptive fields. Based on chapters 4 and 5, a lack of center-surround LGN fields could account for the discrepancies seen in spatial frequency tuning and size tuning for flashing spots. In chapter 5.2, the DoG model [93] was introduced, a model for describing center-surround organisation of receptive fields. Using this model, the LGN spatial filters are modified from center-only to center-surround. The simulations are then repeated to explore the effect of this modification on the LGN and V1. Bear in mind that the derived responses apply to simple circular and center-surround receptive fields, corresponding to ON and OFF cells. A fraction of cells in the LGN model are ON/OFF cells, which have more complex responses. Since there are relatively few of these cells (11.4 %) [12], the focus will be on the response of the ON and OFF cells, although it will be shown that similar results were found for the ON/OFF cells.

One type of experimental finding not fully explained by the center-surround receptive fields of the LGN is size suppression for patch gratings in V1. As explained

in section 4.1.1, the mechanisms behind this suppression are not yet understood, and may include network effects within LGN, V1 or both, as well as feedback and feed-forward signaling. Similarly, there is not yet agreement on how to best model this phenomena [70,125,129]. Some findings suggest strong involvement of horizontal connections in V1 [115]. Size tuning curves of the V1 model with and without recurrent connections are therefore compared, in order to look at the influence of recurrent connections in the model. As an alternative to modelling network effects in a mechanistic way, divisive normalisation has been used by previous studies to represent size suppression in LGN and V1 [70, 129]. To test how this could be used in the Allen model, the LGN were modified so that firing rates are normalised with respect to overall activity. The effect of this modification on size tuning in LGN and V1 is presented.

9.1 The Original Allen Model

9.1.1 Spatial frequency tuning - full-field gratings

LGN

The spatial frequency tuning of two LGN neurons (one ON, one OFF cell) in response to full-field drifting gratings is shown in figure 9.1. The two neurons have receptive field sizes from each end of the size range, defined by equation (5.4), with $\sigma = 3.31°$ and $\sigma = 0.67°$. These parameters give rise to the lowest and highest preferred spatial frequencies of the LGN cells respectively. Tuning curves for an additional two neurons having receptive field sizes defined by $\sigma = 2.45°$ and $\sigma = 1.56°$ are given in the appendix (see figure B.1). The results from simulations (black dots) are compared with the result predicted by equation (5.8) (grey line), in order to confirm that the LGN behaved as expected. This will be particularly important in section 9.2.3, where a key assumption is that the spatial frequency tuning of the LGN neurons can be predicted analytically.

The analysed cells display low-pass spatial frequency tuning. This type of tuning is a property of linearly behaving cells without surround antagonism (see section 4.2 and 5.2), such as the ON and OFF cells in the original LGN model. The other type of LGN neurons, the ON/OFF cells, have receptive fields consisting of two Gaussian filters separated by a few degrees in visual space: one main Gaussian spatial filter and one smaller Gaussian of opposite sign (see section 6.2). Unlike the ON and OFF neurons, their spatial frequency tuning can therefore not be predicted by equation (5.8). Still, all ten analysed cells from each of the ON/OFF groups (sONsOFF, sONtOFF) also display low-pass tuning, as seen in figures B.2 and B.3 in the appendix.

The frequency tuning in figure 9.1a and that of the two ON cells in the appendix (figure B.1) follow the analytical solution. For high spatial frequencies however, the analytical and simulated results diverge (figure 9.1b). This is most likely caused by the inadequate resolution of visual stimuli with high frequencies. Better agreement between the model and analytic solution was found for increasing grating resolutions. The resolution of the visual input was set to 2 pixels per degree, which cannot accurately represent gratings of the highest spatial frequencies, but was chosen as a compromise with computational run-time. It will therefore be assumed that for grating input which is exactly described by equation (5.7), the LGN output follows the analytical tuning curve.

In summary, the ON and OFF cells seem to follow the analytical solution, which is defined by equation (5.2) as being low-pass. Additionally, all analysed ON/OFF cells (a class comprising 11.4% of LGN cells) display low-pass tuning. This suggests

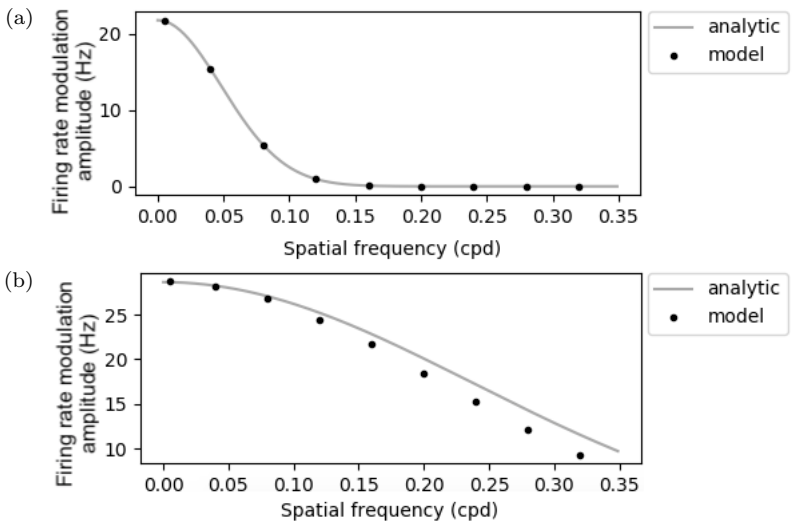

Figure 9.1: Spatial frequency tuning of two neurons in the original Allen LGN model. The F1 response of the analytical solution from equation (5.8) (grey line) and simulated results (black dots) are compared. The neurons have receptive field sizes defined by equation (5.4), with $\sigma = 3.31°$ (a) and $\sigma = 0.67°$ (b). The neuron ids and classes are 17173, sOFF-TF8 (a) and 3099, sON-TF4 (b).

that the LGN model consists of exclusively or nearly exclusively of cells with low-pass spatial frequency tuning. In contrast, experiments on mice LGN have measured band-pass tuning in a large proportion of cells (see section 7.1.1).

V1

Spatial frequency tuning curves were obtained for V1 neurons in the Allen model. The V1 network was driven by the output rates of the LGN, generated in response to full-field drifting gratings of various spatial frequencies. The results for 10 responsive neurons from each of the V1 cells types are shown in figure 9.2. Figure 9.3 presents the optimal frequencies of V1 neurons by type. 50 cells have been sampled from each cell class, and only responsive neurons are included in the analysis. Representing the full V1 network, figure 9.4 shows the preferred frequencies of the responsive cells among 1000 randomly sampled V1 neurons.

The majority of V1 neurons show similar low-pass tuning to that of the LGN cells. 67% of cells (615/908 responsive units) prefer the lowest spatial frequency used for visual input (0.005 cpd). On the other hand, 27% of cells (244/908 responsive units) have a peak in activity at 0.04 cpd (figure 9.4). This preference for frequencies above the lowest presented value was not observed in the LGN. This implies that, in some cases, there is a small change in preferred frequencies from LGN to V1. This can be explained by network effects in V1. Since inhibitory V1 cells, like most neurons, show low-pass tuning, inhibition within the network will be strongest at 0.005 cpd. This may lead to cells having higher activities for frequencies above 0.005 cpd, and hence cause the observed variation in preferred frequencies.

Although a number cells show slight band-pass behaviour, most neurons display low-pass tuning, with a median peak frequency of 0.005 cpd. Contrastingly, experimental studies (see section 7.1.2) found that V1 cells, on average, displayed band-pass

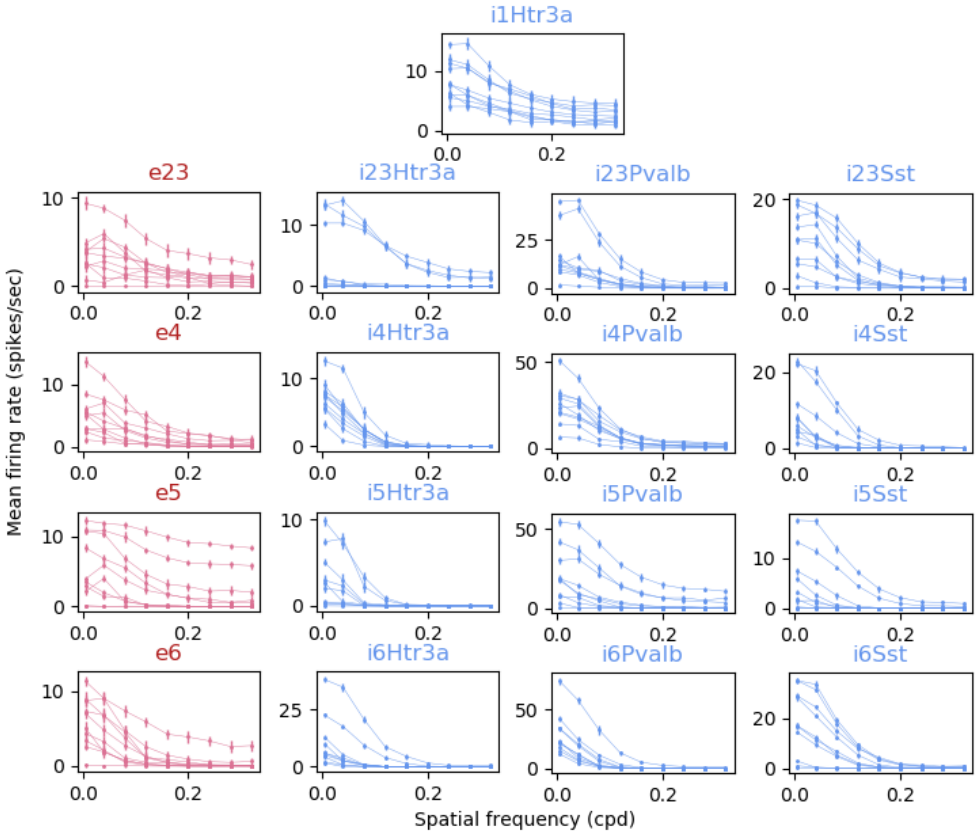

Figure 9.2: Spatial frequency tuning of V1 neurons in the original Allen model. Each subplot shows 10 responsive neurons from the respective classes. Curves show the mean over 20 trials of 3-second simulations, with error bars indicating the standard deviation. The y-axes limits are chosen to show the tuning shape of all the populations. For an overview of how activity vary between populations, look to appendix figure B.6, where the same scale is used for all subplots.

tuning, with median and mean preferred spatial frequencies between 0.04 and 0.5 cpd [16, 116, 118].

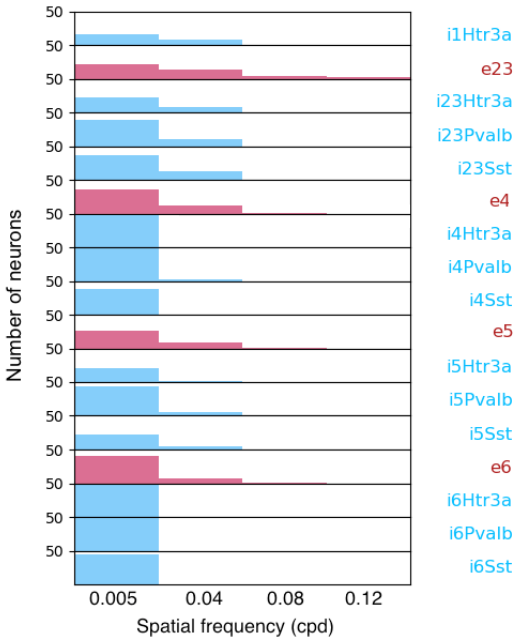

Figure 9.3: Optimal spatial frequencies of V1 neurons in the original Allen model. The labels represent the different V1 classes. For each class, 50 cells are analysed, and the preferred frequencies of the responsive neurons are shown. Gratings of spatial frequencies 0.005, 0.04, 0.08, 0.12, 0.16, 0.20, 0.24, 0.28 and 0.32 cpd are used as visual stimuli. The frequencies 0.16 - 0.32 cpd, to which none of the neurons showed a preference, are omitted from the figure. Preferred frequencies are based of mean firing rates over 20 trials of 3-second simulations.

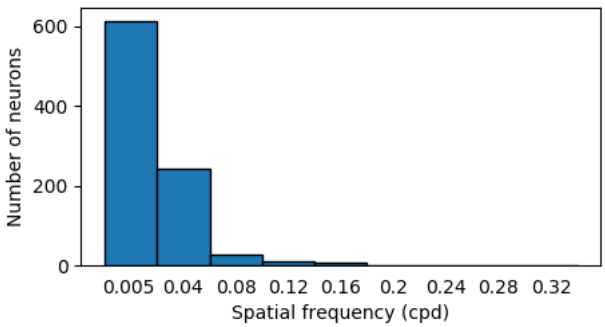

Figure 9.4: Optimal spatial frequencies of the 908 responsive cells out of 1000 randomly sampled V1 neurons from the original Allen model. Firing rates are measured in response to gratings of spatial frequencies 0.005, 0.04, 0.08, 0.12, 0.16, 0.20, 0.24, 0.28 and 0.32 cpd, each frequency represented by a separate bin. Preferred frequencies are based of mean firing rates over 20 trials of 3-second simulations.

9.1.2 Size tuning - flashing spots

LGN

LGN neurons were driven with flashing white spots. Alternating sequences of white spots and grey screen, both of 75 ms intervals, were presented over 3-second simulations. Figure 9.5 shows the activity of (a) eight ON and (b) eight OFF LGN neurons as a function of spot size. These neurons were selected as their receptive field centers

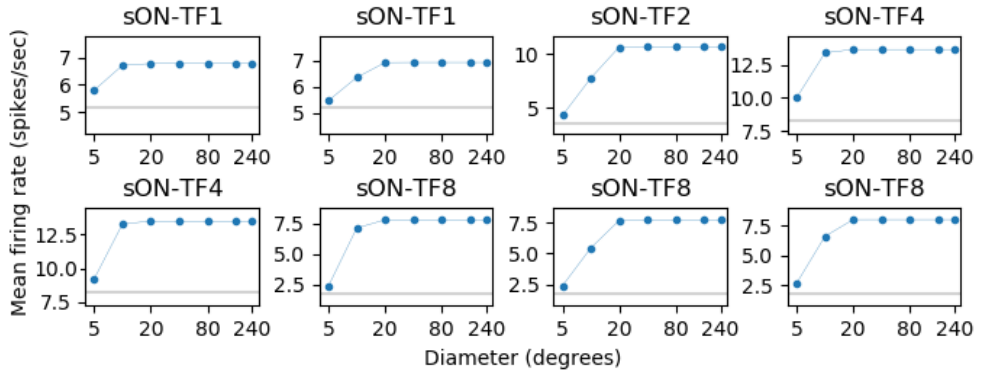

(a)

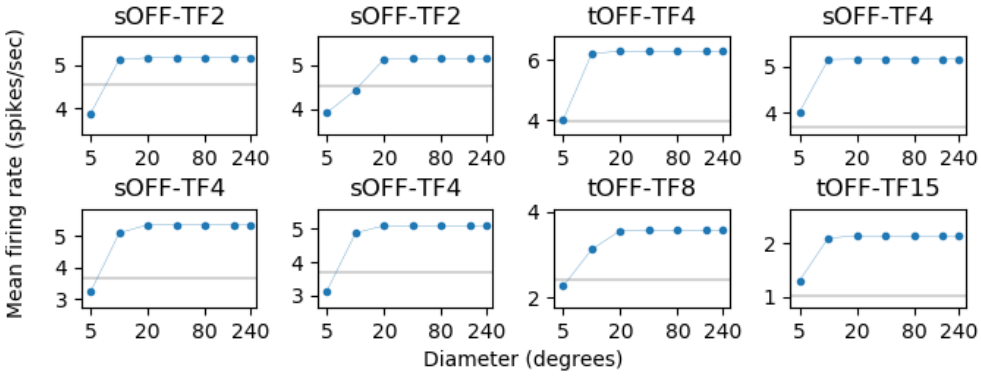

(b)

Figure 9.5: Size tuning (blue line) of LGN ON cells (a) and OFF cells (b). The stimuli was a 3-second sequence of 75ms flashes of a white spot. The spatial LGN filter is the single Gaussian in the original model. Labels indicate the neuron class.

were closest to stimulus center. The rates are the deterministic firing rates generated in the model, rather than recovered from spiking activity. All the recorded LGN cells show the same behaviour for stimuli of increasing sizes: the firing rate increases for diameters up to a certain size, around 10° to 20°, from which point the average activity is constant. This indicates that there is no surround suppression in LGN for flashing spot stimuli.

As explained in section 6.2, the cells in the LGN are modelled as LNP cells with Gaussian receptive fields [12]. Recall that LNP cells act as linear filters coupled with a non-linear rectifier. The rectifying function, given in equation (3.10), sets all negative rates to zero and is linear above zero. The term rectifier is here used with reference this function. The analytic results presented in section 5.2 (see figure 5.3) give the

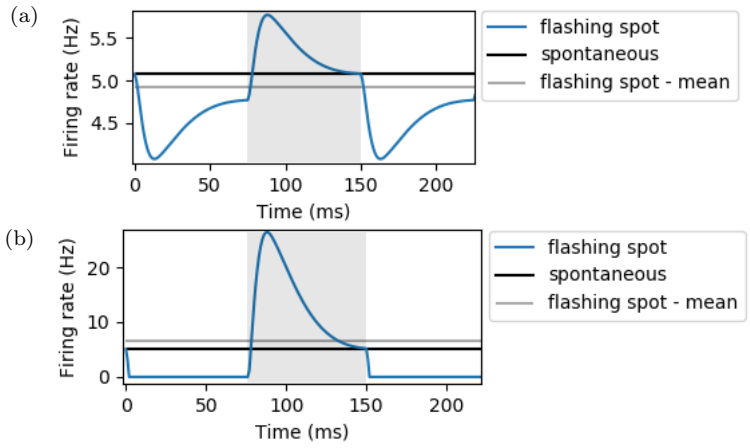

Figure 9.6: Firing rate as a function of time (blue line) and rate average (grey line) for an example OFF neuron (id = 8013) in response to a white flashing spot. The white and grey backgrounds indicate when the white spot is turned on and off respectively. (a) For a spot size of 5° diameter, the momentary firing rate does not reach below the zero-threshold, and the time-averaged response of an OFF-neuron is decreased compared to its spontaneous rate. (b) Spot sizes of 20° diameter evokes a stronger response to both flash onset and offset. Negative firing rates are set to zero by the rectifier, which can cause the time averaged activity to become higher than the spontaneous rate.

linear response of an ON-cell modelled as a single positive Gaussian, and predicts no size suppression for uniform spots of light. The tuning curves in figure 9.5a are consistent with this result.

For an OFF cell, the linear dependence on spot size, as derived in section 5.2, is a tuning curve which is the inverse of the ON cell's. That is, a decrease in firing rate with spot size. This contrasts the tuning curves we see in figure 9.5b. As will be explained below, these seemingly anomalous responses are caused by the rectification of the linear response.

Recall that it is the spatial dependence of mean activity which is predicted in section 5.2, while the response still varies periodically in time. It is therefore the time average of the linear response which follows the shape of the derived tuning curves. The instantaneous response depends on the shape of the temporal receptive field. The temporal kernel in equation (5.5) defines the effect of earlier visual stimuli on LGN activity. The filter changes from positive to negative with increased time since input (see figure 4.8), so that the effect of the current stimulus and that of some milliseconds earlier is opposite [50]. For an LGN OFF-unit, the maximum response will be evoked by stimuli which switches from bright to dark. A spot of high luminosity being switched off can therefore momentarily increase the firing rate in OFF-cells to above the spontaneous rate, as shown in 9.6. For a purely linear OFF-cell, the average activity will still be decreased by the white flashing spots, due to the dip in activity during the flash onset. However, the LGN cells have nonlinear rectifiers preventing the firing rate from becoming negative. This restricts the decrease in firing rate during flash onset, and allows the time averaged activity to increase above the spontaneous rate.

The form of the temporal kernel, combined with the non-linear rectifier, influences the size tuning curve of the OFF-neurons in response to white flashing spots. In

accordance with the linear prediction, the activity initially decreases as a function of stimulus diameter. An illustration of this is the example neuron in figure 9.6, which is inhibited by a white flashing spot of 5° diameter (figure 9.6a). For a certain size, the momentary response to the flash onset reaches the zero-threshold. For stimuli larger than this, the activity will start increasing with stimulus size. For a spot size of 20° diameter, the time averaged activity is above the spontaneous rate (figure 9.6b). The same dependence on stimulus size is seen in the size tuning results of the OFF cells in figure 9.5.

Consistent with the Allen LGN cells, Piscopo et al. [52] found experimentally that both ON and OFF cells may be excited by white flashing spots, although ON cells respond more strongly. However, cell recordings summarised in section 7.1.1 found size suppression for a large proportion of cells in response to flashing spots, something not reproduced by the Allen LGN cells.

V1

For each flashing spot size, the output of the 20 LGN trials were used to drive the V1 model. Size tuning curves were then obtained for single V1 neurons. Up to ten neurons from each population are plotted, depending on the number of cells within a 70 µm radius of the stimulus center (see section 8.7). Figure 9.7 shows the resulting tuning curves, averaged over the 20 trials. The varying degrees of activity between different neurons is influences by several factors, including the amount of input received from LGN and other V1 cells, as well as model parameters such as the spontaneous rate. The activity increases before reaching a plateau for stimuli of around 20° to 40° diameter. This implies that there is no surround suppression in V1 in response to flashing spots.

For comparison, one study in section 7 used flashing spots to map receptive fields in the mouse V1 [117]. The study found that out of the cells with circularly symmetric fields (55%), over half had surround regions which opposed the response of the center [117]. In addition, for the type referred to as simple cells (17%), most cells showed size inhibition as stimuli expanded past the elliptical ON region into the elliptical OFF-regions [117]. Importantly, all cells responded in general either weakly or not at all to full-field changes in luminosity [117]. This is in contrast to the results for the Allen V1 cells, where all responsive units fire optimally for full-field flashes.

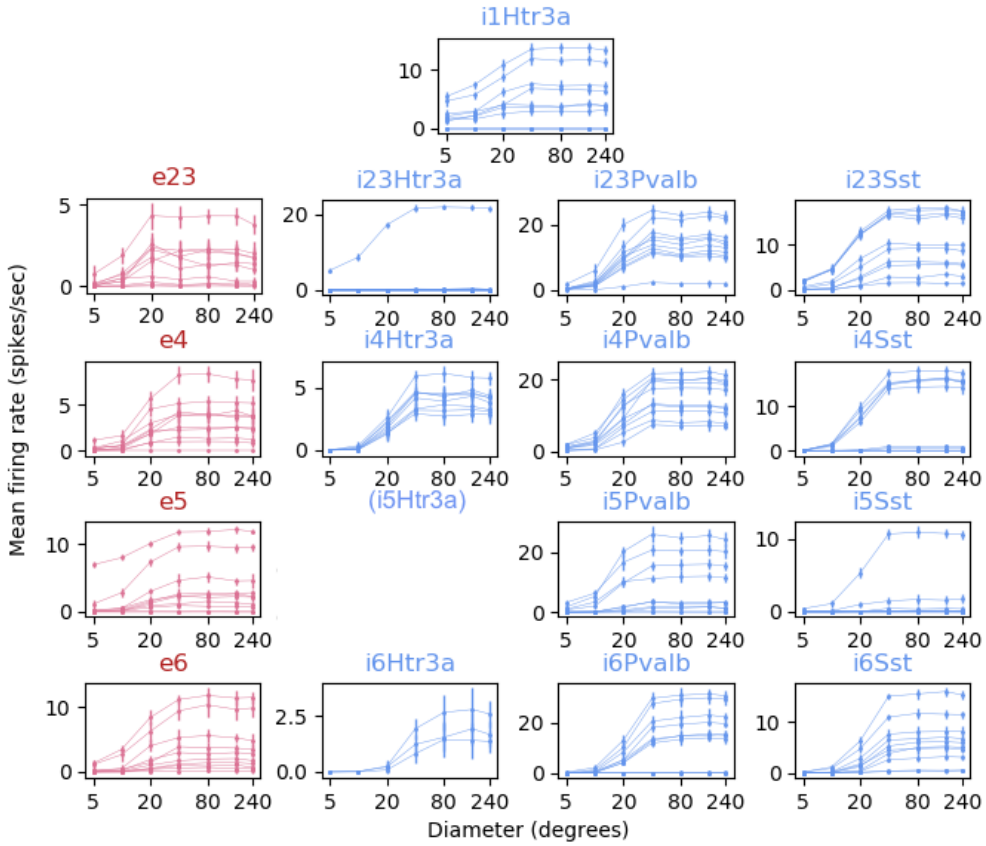

Figure 9.7: Size tuning in V1 for white flashing spots of 75 ms duration. Up to ten neurons from each population are included, depending on the number of cells within 70 µm of stimuli center. For the i5Htr3a class, there were no cells within this radius. The spatial LGN filter is the single Gaussian in the original model. The measurements show mean response over 20 trials of 3-second simulations, with error bars representing the standard deviation. The y-axes limits are chosen to show the tuning shape of all the populations. For an overview of how degree of activity vary between populations, look to appendix figure B.7, where the same scale is used for all subplots.

9.1.3 Size tuning - patch gratings

The second type of stimulus used for size tuning measurements was patch gratings. Initially, simulations with static patch gratings were tested. No measurable response was detected in V1 for averages of up to 60 trials, neither for the standard grating nor for one with an inverted greyscale. Drifting stimuli were therefore used instead, which generated a considerably larger response.

LGN

The LGN model was driven with circular patches of drifting grating over a 3-second simulation. As with flashing spots, different sizes of patch grating were presented, and activity was plotted as a function of stimuli size. Figure 9.8 shows the size tuning of LGN neurons, one from each cell class. The rates are based directly on the deterministic firing rates generated by the LGN model, rather than recovered from spiking activity. Their responses are similar to those for flashing spots, showing no

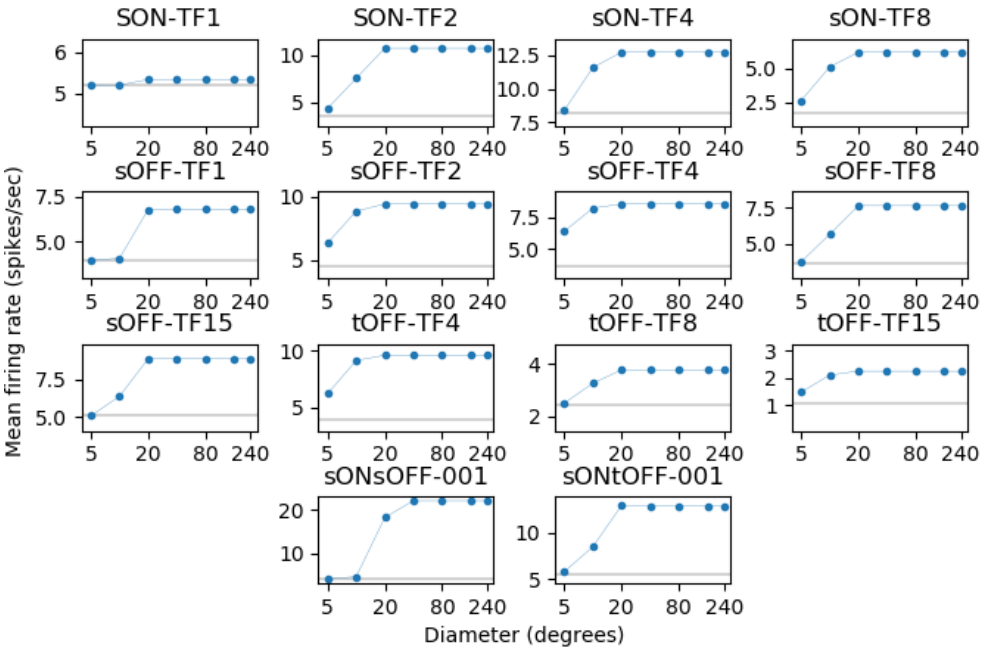

Figure 9.8: Size tuning of LGN neurons for patch grating stimuli. The tuning curve of one example neuron from each cell class is plotted. The LGN spatial filters used are the single Gaussians of the original model.

size suppression for any of the recorded cells. As before, this can be explained by the cells being linear filters (except for a nonlinear rectification), with no surround antagonism to oppose the center.

V1

For size tuning with patch gratings, the spike trains from 20 LGN trials were used to drive the V1 model. The results are shown in figure 9.9. The figure presents the activity of up to ten V1 neurons from each population, depending on the number of cells within a 70 µm radius of the cortex center. The activity represent the time

average over 20 trials. The neurons are the same as those measured for flashing spots. As for flashing spots, the V1 neurons do not display surround suppression for patch gratings. This is in contrast to the experimental findings of size suppression in V1 in response to patch gratings (section 7.1.2).

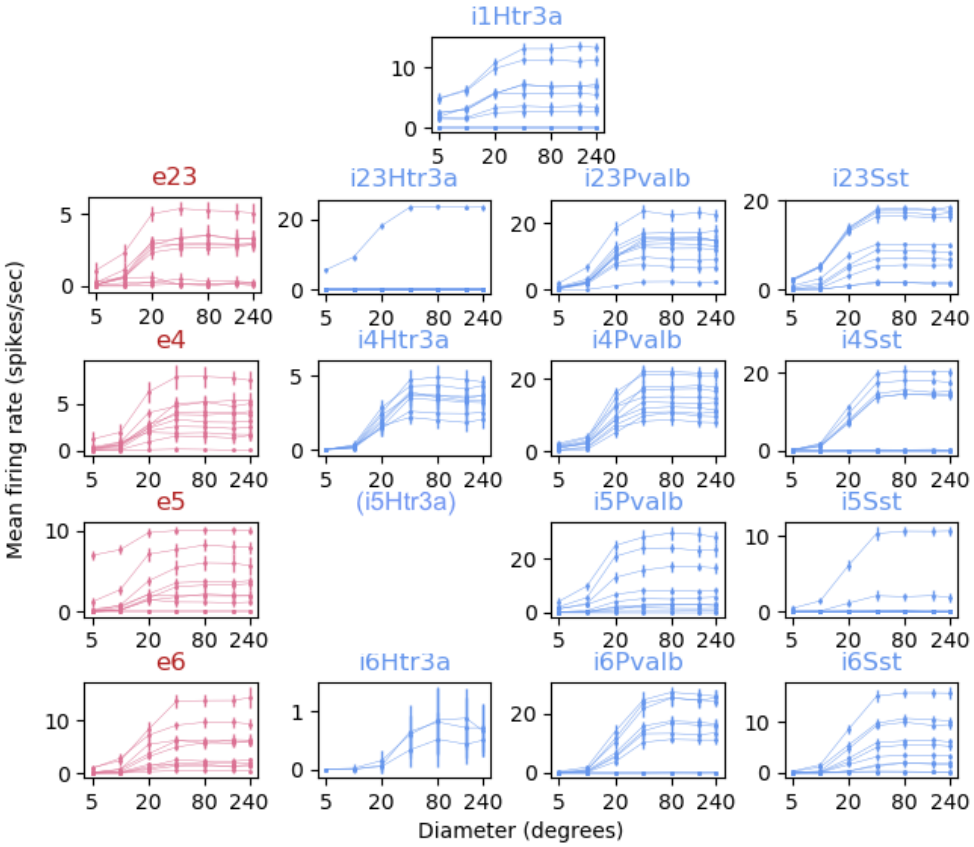

Figure 9.9: Size tuning of V1 neurons for patch grating stimuli in the original Allen model. Up to ten neurons from each population are included, depending on the number of cells within 70 µm of stimuli center. For the i5Htr3a class, there were no cells within this radius. Mean firing rates over 20 trials of 3-second simulations are plotted, with error bars showing the standard deviation. The y-axes limits are chosen to show the tuning shape of all the populations. For an overview of how degree of activity vary between populations, look to appendix figure B.10, for which the same scale is used for all subplots.

9.2 The M1 Allen Model

The M1 model refers to a modified version of the Allen model. In particular, the single Gaussian spatial filters of the original LGN have been replaced with DoG filters, defined by equation (5.6), with $A_c = A_s = 1.73$ and $\sigma_s = 2.45\sigma_c$. The following section explains the motivation for modifying the model. Analytical results used for tuning the DoG parameters are then presented, before moving on to LGN and V1 simulation results.

9.2.1 Motivation for modifying the model

The results for the original Allen LGN model show no surround suppression for flashing spots or patch gratings. In addition, all the LGN cell types display low-pass tuning, preferring the lowest spatial frequency presented (0.005 cpd). Each LGN cell consists of a linear filter coupled with a non-linear rectifier, as defined by equation (3.10). The low-pass tuning and lack of surround inhibition in the model can be explained by the cells having only an excitatory or inhibitory center region, so that there is no on-off antagonism that can give rise to band-pass tuning or size suppression, neither for flashing spots nor patch grating stimuli.

While the Allen LGN is an LNP filter model, the Allen V1 is a nonlinear model with multiple recurrent connections, where neurons receive input from both LGN neurons and other V1 neurons [12]. In real mice, the recurrent connections in V1 have been suggested to be at least partly responsible for nonlinear surround suppression [74]. However, none of the analysed V1 cells show surround suppression for flashing spots or patch grating stimuli. While a substantial proportion of V1 cells prefer gratings with spatial frequency of 0.4 cpd, most cells prefer the lowest frequency of 0.005 cpd.

In contrast, experimental studies found surround suppression in LGN in response to flashing spots [52,83,113], and in V1 in response to patch gratings [48,64,82,115,116] and flashing spots [117] (see chapter 7). For experiments measuring spatial frequency preferences, band-pass tuning was found in the majority of LGN [16, 28, 52, 83, 113] and V1 cells [16, 116, 118]. The further aim was therefore to modify the Allen LGN so that the full network model will exhibit these properties.

In section 5.3, analytical solutions for the spatial-filter response to bright spots and drifting gratings are calculated. It is explained how the DoG function is able to account for surround suppression in response to uniform bright stimuli and for the spatial frequency tuning curves which are found experimentally in mice. In some cases, the DoG filter also exhibits surround suppression with patch gratings, although this requires specific combinations of spatial frequency of stimuli and filter parameters.

To reproduce band-pass tuning and surround suppression in the Allen LGN, the spatial filter was modified from the single Gaussian in equation (5.6) to the DoG function in equation (5.4). The following chapter presents the resultant LGN tuning properties. The effect of the modification on V1 is also explored, by repeating the same measurements as when using input from the original LGN model.

9.2.2 The DoG filter

The method for modifying the Allen LGN model and choosing parameters is described in section 8.1.1. In short, the equation of the spatial filter was changed from 5.4 to 5.6, with the following parameters.

- $A_s = A_c = 1.73$. The values of A_s and A_c are set equal in order to obtain maximal surround inhibition, as the aim was to see clearly which properties were displayed by the model. For a weaker, and possibly more realistic suppressive effect (see experimental findings 7.1.1) one might want to use a smaller amplitude of the surround Gaussian. Both amplitudes have been scaled (see section 8.1.1) so that the peak firing rate of the DoG spatial frequency tuning curve matches the corresponding peak rate of the original Gaussian filter.

- σ_c for each cell is equal to the value of σ which was used in the original model. The original receptive field sizes in the LGN model are based on recordings by Durand et al. [16], and are similar to the sizes of the ON and OFF areas reported in this study.

- σ_s is defined by $\sigma_s = C\sigma_c$, where $C = 2.45$ has been tuned (see section 8.1.1) so that the median preferred frequency is 0.08 cpd, the same as found empirically by Durand et al. [16]. This will be described in further detail in section 9.2.3 below.

9.2.3 Analytical parameter tuning: peak LGN frequencies

Section 4.2 explains the relationship between size tuning and spatial frequency tuning in cells with a high degree of linearity. Section 5.3.1 describes mathematically how the center-surround organisation which explains surround suppression for uniform stimuli also gives rise to band-pass spatial frequency tuning. The value of the peak frequency depends on the size of the center and surround of the receptive field. The relationship between spatial frequency tuning and receptive field structure was used to determine the size of the inhibitory Gaussian in the DoG filter.

Specifically, equation (5.9) was used to find the optimal spatial frequencies of all the LGN cells. This rests on the assumption that the cells' spatial frequency tuning follow the analytical solution, as supported by findings in the original LGN (see section 9.1.1 and appendix B.1). The parameter defining the size of the inhibitory Gaussian, $C = 2.45$, was then set to give a median optimal frequency of 0.08 cpd across all LGN cells (see method in section 8.1.1). This is the median value which was found in real awake mice LGN by [16].

The resulting histogram of preferred spatial frequencies for all DoG LGN neurons is shown in figure 9.10. As an approximation, all the LGN neurons are considered as simple ON or OFF cells. This includes the ON/OFF neurons (11.4% of all cells [12]), which are treated as having a single, circular receptive field, as opposed to one large and one small circular region. The left and right sharp peaks belong to the sONtOFF and sONsOFF cells, for which σ_c is equal to 3 and 2 degrees respectively. The range of preferred spatial frequencies were 0.041 to 0.201 cpd, which is narrower than the range from 0.02 to 0.64, as found experimentally by Durand et al. [16]. To approach the range measured in [16], one could have increased the variation of peak frequencies by letting C in $\sigma_s = C\sigma_c$ vary between cells. One option is to draw C from some distribution. The distribution could then have been adjusted to obtain a range more similar to that in [16]. However, for now the primary aim was to see which phenomena could be reproduced in the resulting modified model, rather than fine tuning all the parameters.

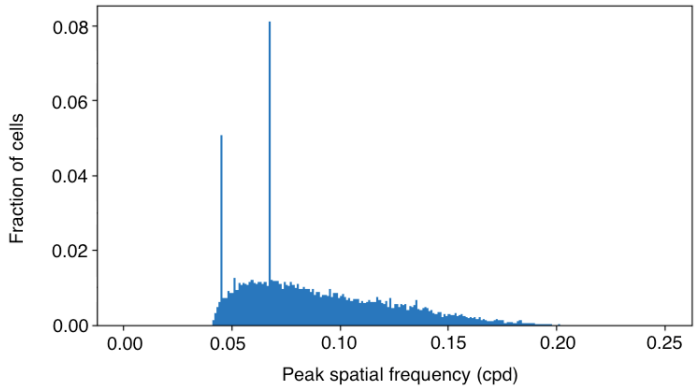

Figure 9.10: Histogram of predicted preferred spatial frequencies for all the 17 400 neurons in the LGN, based on equation (5.9), where $\sigma_s = 2.45\sigma_c$. The width of each bin is 0.001 cpd. The median preferred spatial frequency is 0.08 cpd. The two large peaks are sONtOFF (left) and sONsOFF (right) cells, for which σ_c is equal to 3 and 2 degrees respectively.

9.2.4 Spatial frequency tuning - full-field gratings

LGN

An important assumption when calculating peak spatial frequencies in LGN, and thus determining C in $\sigma_s = C\sigma_c$, was that the spatial frequency tuning curves could be predicted analytically. The spatial frequency tuning of the LGN cells in response to full-field drifting gratings was therefore compared with the analytical result for the selected DoG filter. In this way, the analytical solution served as a benchmark, ensuring that the modified DoG model behaved as expected. Figure 9.11 shows the comparison for an ON and an OFF cell. The cells have receptive field sizes from each end of the size range, with $\sigma_c = 3.31^\circ$ and $\sigma_c = 0.67^\circ$, giving rise to the lowest and highest preferred spatial frequencies for the LGN neurons respectively. Curves for two additional cells with receptive field sizes defined by $\sigma_c = 2.44^\circ$ and $\sigma_c = 1.56^\circ$ are given in the appendix (figure B.11).

The simulation output overlaps the analytical prediction for the cell in figure 9.11a and for the two neurons in the appendix (figure B.11). For the cell preferring high spatial frequencies, the analytical and simulated results diverge (figure 9.11b). The same was found for the original Allen LGN (see section 9.1.1 and appendix B.1). This could be explained by a grating resolution too low to resolve high spatial frequencies (see section 9.1.1). It is therefore assumed that for visual inputs which are exactly a sinusoidal grating, the ON and OFF cells reproduce the analytical results given by equation (5.9).

The simulated tuning curves of the ON/OFF cells are presented in figure B.12 and B.13 in the appendix. For the majority of the analysed ON/OFF cells, their preferred frequencies were similar to the analytical prediction, where the neurons were approximated as simple ON or OFF cells. This implies that even with this approximation, the predicted peak frequencies for the ON/OFF cells in figure 9.10 correspond well with the simulations.

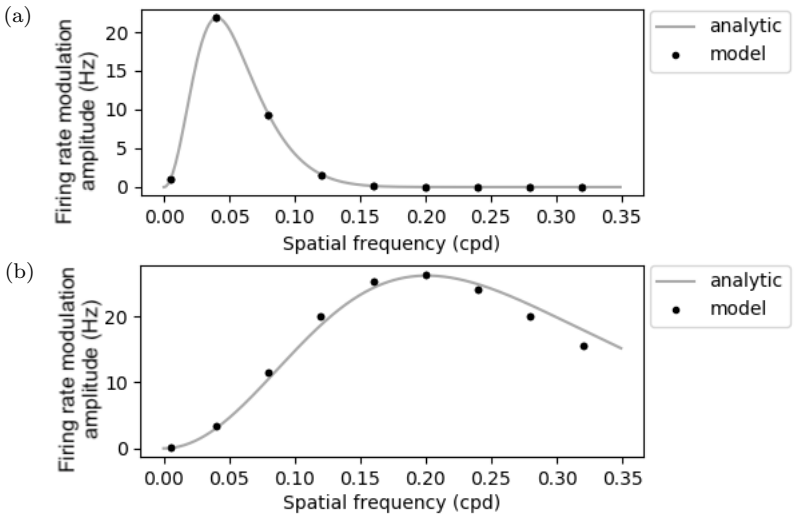

Figure 9.11: Spatial frequency tuning for two neurons from the modified LGN model. The analytical solution from equation (5.9) (grey line) and simulated results (black dots) are compared. The neurons have spatial receptive fields defined by equation (5.6). (a) The neuron (id = 17173, type = sOFF-TF8) has receptive field size with $\sigma_c = 3.31$. (b) The neuron (id = 3099, type = sON-TF4) has a receptive field size with $\sigma_c = 0.67$.

V1

Spatial frequency tuning curves for V1 were found from model simulations, using the same method as for the original Allen model. Tuning curves for 10 neurons from each of the V1 cell classes are shown in figure 9.12. Out of 50 cells sampled from each class, the preferred frequencies of the responsive neurons are shown in figure 9.13. To represent the total V1 network, the peak spatial frequencies of the responsive neurons out of a 1000 randomly chosen cells are shown in figure 9.14.

All recorded cells display band-pass tuning. The median preferred frequency in V1 is 0.08 cpd, the same as for LGN neurons. Durand et al. [16] found a slightly lower optimal frequency (median = 0.05 cpd) in V1 of awake mice (see chapter 7). However, the difference between this value and their LGN results (median = 0.08 cpd) was not significant, indicating no discrepancy with the Allen results. Nevertheless, similar results in other studies (mean = 0.04 [116] and 0.045 [118]) indicate that the preferred spatial frequency may be slightly lower in real mice neurons than in the Allen neurons.

The histogram in figure 9.14 is quite narrow, with 73% of V1 cells (651/889 responsive cells) preferring the grating of 0.08 cpd. This variation is much less than that seen in LGN. A possible reason could be that input from LGN cells average out, so that V1 cells receive the most input for 0.08 cpd, the median preferred frequency of the LGN. By contrast, the cells studied by Durand et al. [16] had a broad distribution of peak frequencies, with a large proportion of cells firing in the range $0.02 - 0.32$ cpd. Similarly, the study by Marshel et al. [118] (see chapter 7) showed a high variety of optimal frequencies, with cells displaying both low-pass, high-pass and band-pass behaviour. In other words, more variation was found in the real mouse V1 neurons than in the simulated Allen V1 neurons in terms of preference for spatial frequency.

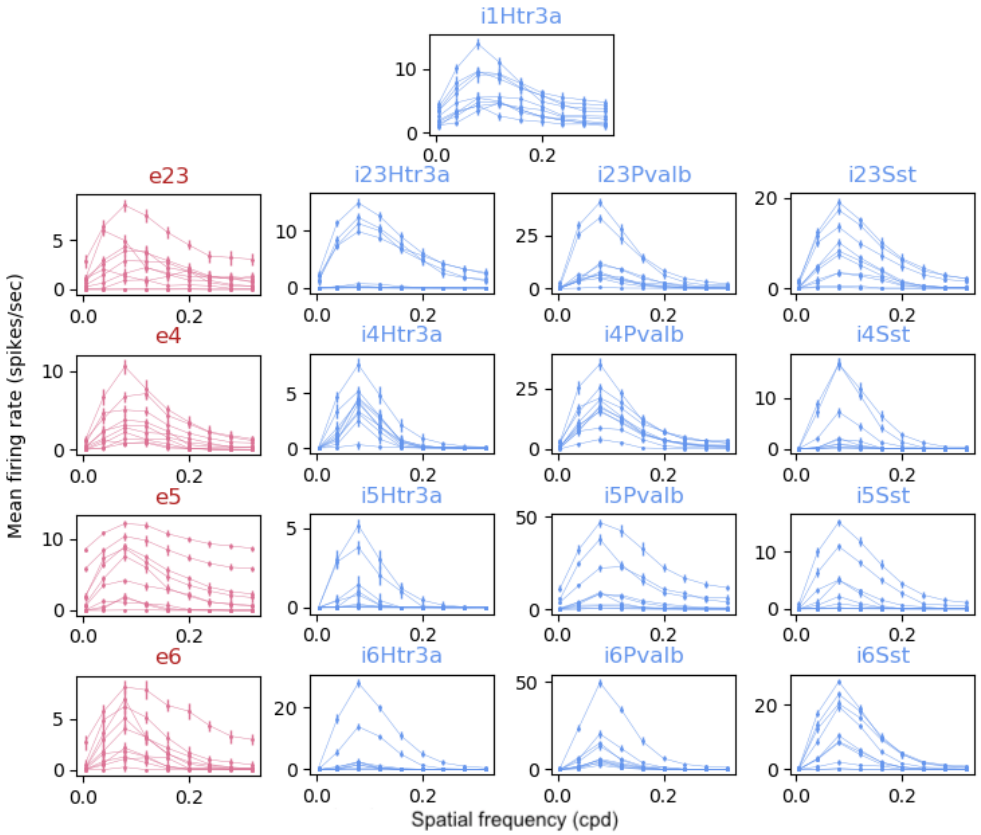

Figure 9.12: Spatial frequency tuning of V1 neurons in the M1 Allen model. Each subplot shows 10 responsive neurons from the respective classes. Curves show the mean over 20 trials of 3-second simulations, with error bars indicating the standard deviation. The y-axes limits are chosen to show the tuning shape of all the populations. For an overview of how the degree of activity vary between populations, look to appendix figure B.16, where the same scale is used for all subplots

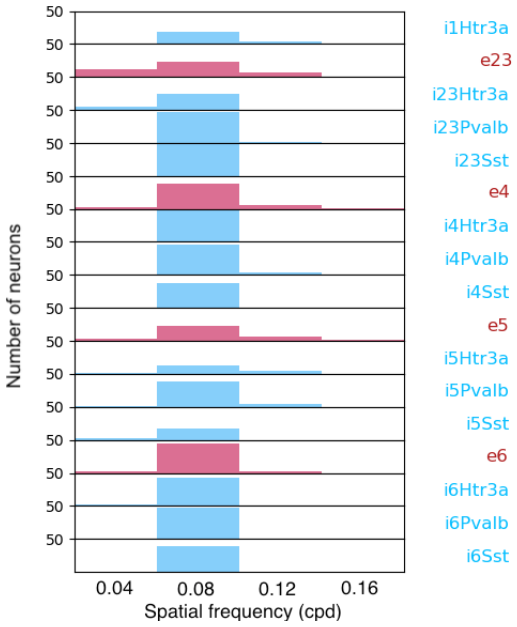

Figure 9.13: Optimal spatial frequencies of V1 neurons in the M1 Allen model. The labels represents the different V1 classes. For each class, 50 cells are analysed, and the preferred frequencies of the responsive neurons are shown. Gratings of spatial frequencies 0.005, 0.04, 0.08, 0.12, 0.16, 0.20, 0.24, 0.28 and 0.32 cpd are used as visual stimuli. The frequencies 0.005 cpd and 0.20 - 0.32 cpd, to which none of the neurons showed a preference, are omitted from the figure. Preferred frequencies are based of mean firing rates over 20 trials of 3-second simulations.

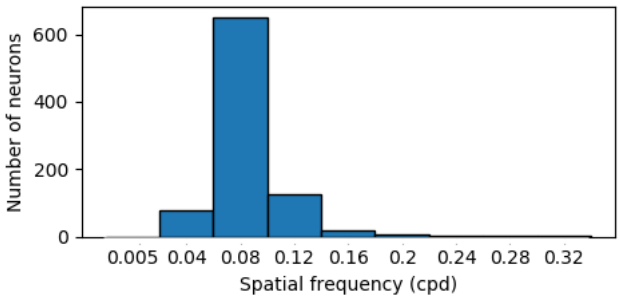

Figure 9.14: Optimal spatial frequencies of 889 responsive cells out of 1000 randomly drawn neurons in V1 for the M1 Allen model. The activity was measured in response to gratings of spatial frequencies 0.005, 0.04, 0.08, 0.12, 0.16, 0.20, 0.24, 0.28 and 0.32 cpd, each frequency represented by a separate bin. Preferred frequencies are based of mean firing rates over 20 trials of 3-second simulations.

9.2.5 Size tuning - flashing spots

LGN

LGN neurons were driven with an alternating sequence of white spots and grey screen, each of 75 ms intervals, over 3-second simulations. Figure 9.15 shows the activity of eight M1 ON and OFF neurons in the LGN as a function of spot size. These are the

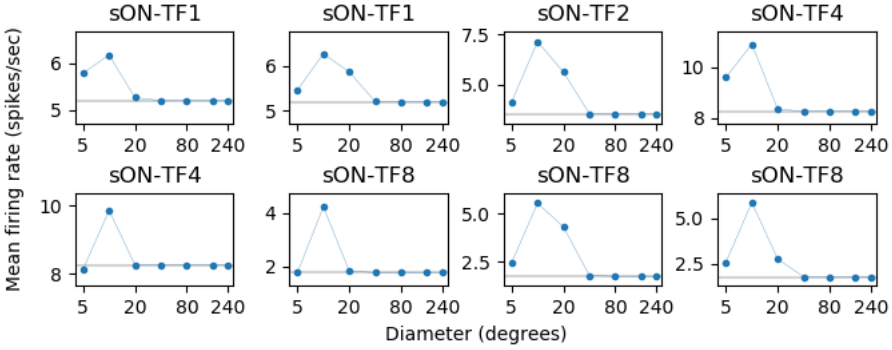

(a)

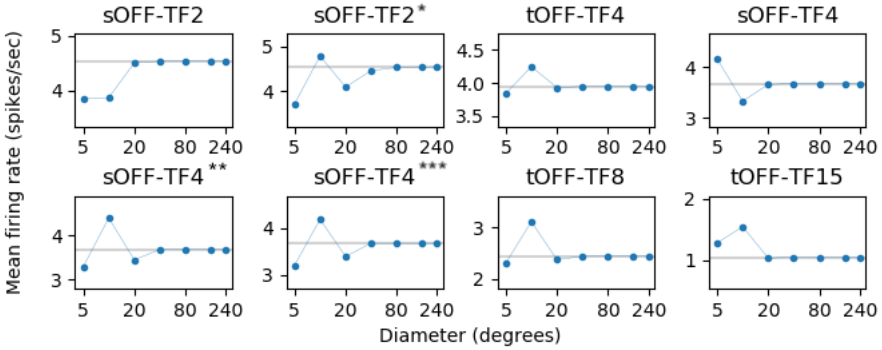

(b)

Figure 9.15: Size tuning for eight ON (a) and OFF (b) cells in the M1 LGN. Flashes of white spots with 75 ms duration were used as stimuli to drive the neurons. Blue dots show measured average firing rate for the different spot sizes, while the grey line marks the spontaneous activity of the neuron. The marks $*$, $**$ and $***$ are for use in the main text when referring to subfigures.

ON and OFF neurons with visual centers closest to the spot center. The rates are the deterministic firing rates generated by the model, rather than rates recovered from spiking activity.

Figure 9.15a shows that surround suppression is displayed by ON cells in the M1 LGN model. This agrees with the analytical linear response of section 5.3.2, which predicts size suppression for on-center/off-surround receptive fields in response to bright stimuli. The results are also consistent with empirical observations in mice [16, 28, 52, 83, 113]. The largest firing rate in LGN was measured for spots of 10° diameter, which implies that the preferred stimuli spot diameter for most ON cells is around 10°. Similarly, Tschetter et al. [83] found that LGN cells had a mean preferred diameter of 8.6°, and Piscopo et al. [52] reported a mean optimal diameter of 11.6° for sON LGN cells, the same type of cells which are analysed here.

Figure 9.15b shows surround suppression for most of the recorded OFF cells in the M1 LGN. The size tuning of the cells may seem to contradict the response predicted analytically in section 5.2. In section 5.2, the size tuning of an on-center cell for bright stimuli was described. In the case of an off-center cell, the linear dependence on stimulus size is the inverse of the DoG response in figure 5.3. However, as explained in section 9.1.2, the nonlinear rectifier of the OFF cells changes the shape of their size tuning. Section 9.1.2 also describes how this depends on the temporal kernel.

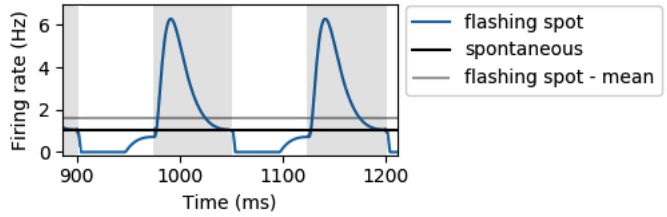

(a)

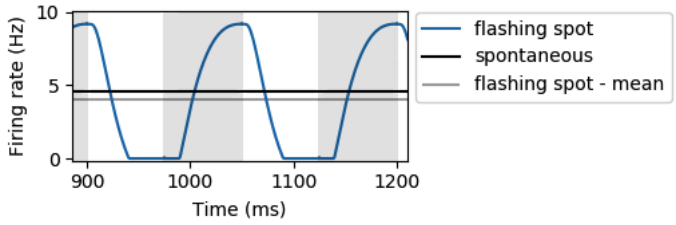

(b)

Figure 9.16: Firing rate average (grey line) and as a function of time (blue line) for M1 OFF neurons in response to white flashing spots of 10° diameter. The black line marks the spontaneous firing rate. The time window is chosen for a time where the amplitude has stabilised. The background colour indicates the alternation between white spots and grey screen. (a) The neuron has a low spontaneous rate compared to the peak in activity, which leads to a time average which is higher than the spontaneous rate. This leads to surround suppression (bottom right, figure 9.15b). (b) The neuron has a high spontaneous rate compared to the peak in activity, so that the time average is lower than the spontaneous rate. This does not lead to surround suppression (top left, figure 9.15b).

The size tuning of the M1 OFF cells can be explained by referring to the activity of neurons (marked $*$, $**$, $***$) in figure 9.15b. That is, initially, the response is linear, and the activity decreases below the spontaneous rate (5° diameter). As soon as the inhibition by the flash onset causes the momentary firing rate to dip below zero, the threshold comes into effect, preventing negative firing rates. Since the excitatory effect of flash termination (the "rebound") is unaffected, this causes the firing rate to start increasing as a function of spot diameter (10° diameter). As the stimuli diameter increases into the surround area, the on-off contributions start cancelling out, and the rate will decrease again (20° diameter). When the momentary inhibitory response to flash onset does no longer reach the zero-threshold, the response becomes linear again and increases until the on-off contributions have completely cancelled, giving the spontaneous firing rate (40° diameter).

If the spontaneous rate is sufficiently low compared to the excitatory rebound, the average activity in the M1 OFF-cells can be increased above the spontaneous rate. This is illustrated by the example neuron in figure 9.16a. Then, since the activity

for large enough stimuli is just the spontaneous rate, this neuron displays surround suppression (bottom right, figure 9.15b). In figure 9.16b on the other hand, the neuron has a relatively high spontaneous rate compared to the rebound from the flash offset. The average activity is therefore either lower or equal to the spontaneous rate for all measured diameters, so that no surround suppression is seen (top left, figure 9.15b).

To summarise, all M1 ON neurons show surround suppression for white flashing spots. For OFF cells, the response is much weaker, and some show surround suppression, while some do not. To see if the result was affected by spot duration, the measurements were repeated for white flashing spots of 300 ms duration (appendix figure B.17 and B.18). The same results were found in this case, with all ON cells and some OFF cells showing suppression. Correspondingly, the studies on mouse LGN in chapter 7 reported a large proportion of cells with surround suppression for flashing spots. Similar to in M1 LGN cells, Piscopo et al. [52] also reported that OFF cells could be excited by white flashing spots, but that the response was weaker than that of ON cells.

V1

The LGN responses to the 75 ms flashing spots were used to drive the V1 network. For each spot size, the 20 LGN trials were used to create 20 corresponding spike trains in V1. The responses of same V1 neurons as previously were then analysed, and their size tuning plotted in figure 9.17. In addition to the 75 ms flashes, measurements were repeated for 300 ms flashes (see figure B.20 in the appendix), which gave the same type of size tuning.

Most of the neurons show size suppression for flashing spot stimuli, a property also reported by Dräger et al. [117]. Furthermore, according to spatial frequency analysis (section 4.2), size suppression in response to flashing spots is consistent with band-pass spatial frequency tuning, which was also found experimentally in mice V1 [16, 116, 118]. The cells show different preferred diameters, lying somewhere between 10° and 40°. In comparison, the experimental literature gave very varying results, ranging all the way from 17.5° (median) to 86° (mean), depending on cell type and degree of anaesthesia [64, 82, 115, 116].

The peak responses are fairly weak compared to those of the original model. This may be because the response of each LGN neuron which transmits to a V1 cell begins to diminish after stimuli exceed the center of the LGN's receptive field. To get a stronger response in V1 cells, the surround amplitude of the DoG filter could be decreased compared to the center amplitude. This would make the surround effect weaker, and probably also more realistic, as the experimental mean suppression index reported in LGN was 0.53 [83]. (Recall section 4.1.2 for an explanation of suppression index.) The M1 model, however, has equal values of A_c and A_s for the center and surround Gaussians in equation (5.6). This maximises surround suppression in response to uniform stimuli, giving SI values of 1 in LGN and decreased input to V1.

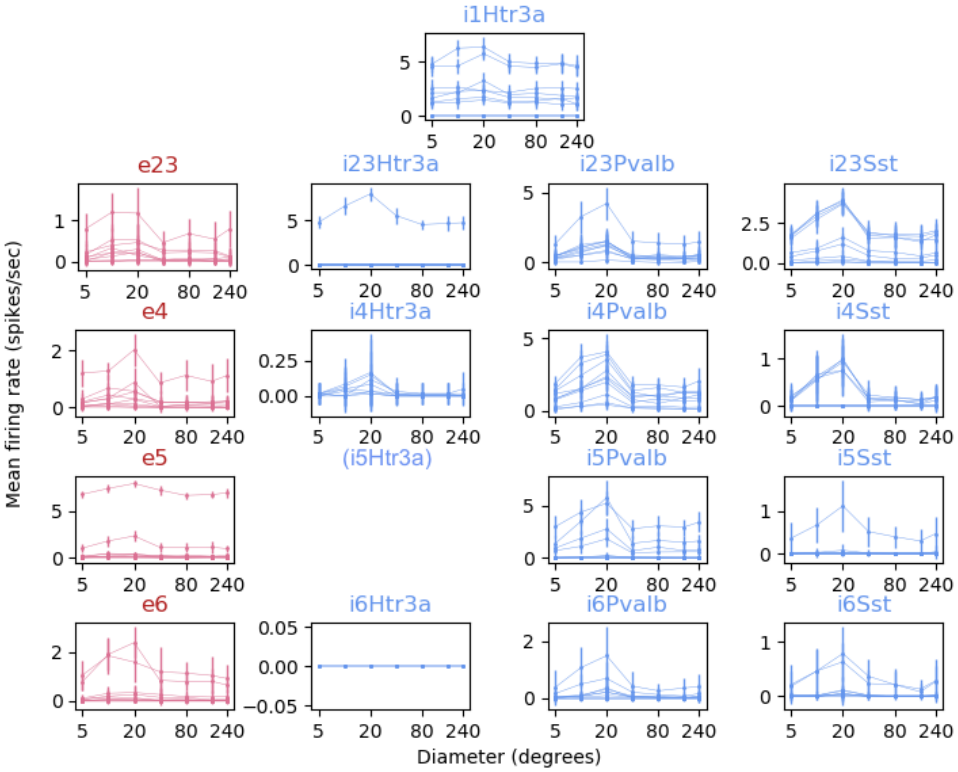

Figure 9.17: Size tuning of V1 neurons from the M1 Allen model, using stimuli consisting of 75 ms flashes of white spots. Up to ten neurons from each population are included, depending on the number of cells within 70 µm of stimulus center. For the i5Htr3a class, there were no cells within this radius. The measurements show mean response over 20 trials of 3-second simulations, with error bars representing the standard deviation. The different y-axes limits are chosen to show the tuning shape of all the populations. For an overview of how the degree of activity vary between populations, look to appendix figure B.19, for which the same scale is used for all subplots.

9.2.6 Size tuning - patch gratings

LGN

Patches of drifting gratings were presented as stimuli to the LGN model, and activity as a function of stimuli size was measured. Simulations were run for gratings with spatial frequencies of 0.02, 0.04, 0.08, 0.12, 0.16 and 0.20 cpd. Figure 9.18 shows the resulting size tuning of M1 LGN neurons for patch grating stimuli of frequency 0.04 cpd. Tuning curves for the remaining frequencies are given in figures B.21 (0.02 cpd), B.22 (0.08 cpd), B.23 (0.12 cpd), B.24 (0.16 cpd) and B.25 (0.20 cpd) in the appendix. The tuning curves in figure 9.18 show a variety of responses, where some

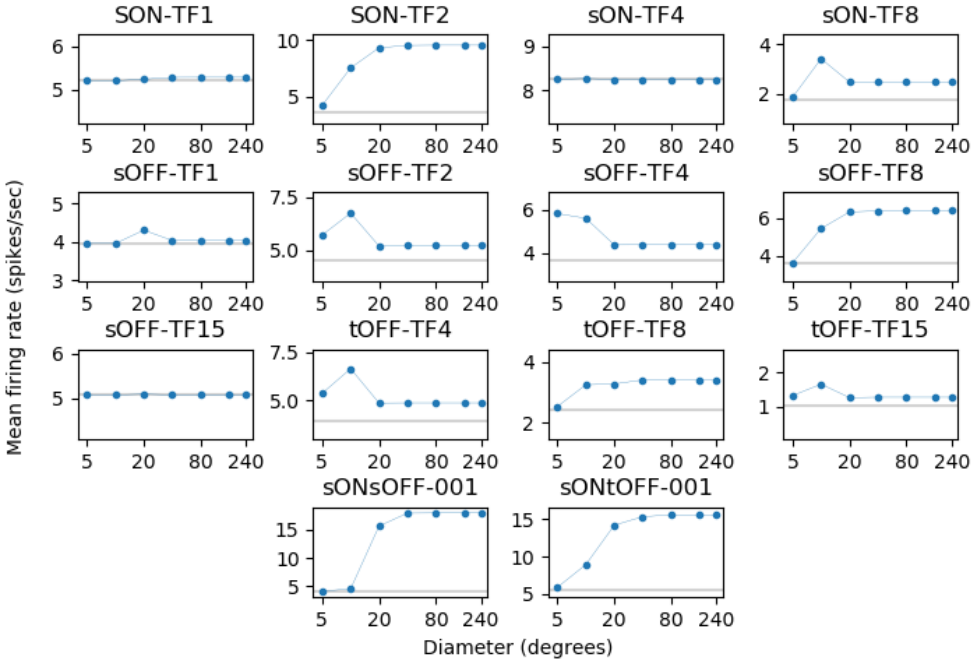

Figure 9.18: Size tuning for one example neuron of each class in the M1 LGN. The stimuli is a patch grating of spatial frequency 0.04 cpd. The spontaneous rates of the cells are marked by the horizontal grey lines.

show surround suppression and some do not. Similarly, the results for the remaining frequencies included cells both with and without surround suppression.

In section 5.3.3, the response of DoG filters to patch gratings is described. In particular, we see that size tuning depends on the parameters σ_c and σ_s, as well as the grating's spatial frequency. These results can be used to explain the tuning curves in figure 9.18. The input to all the LGN neurons has the same spatial frequency, that of the grating, while each σ_c and σ_s differ. For each cell, $\sigma_s = 2.45\sigma_c$, where σ_c is drawn from a distribution centred around the experimentally observed value. It is therefore expected that each neuron has a distinct size tuning curve. As predicted in section 5.3.3, the size tuning for low spatial frequencies (0.02 cpd) is similar to that for flashing spots. For this frequency, all responsive neurons showed suppression, however most neurons responded weakly or not at all. This suppression can be understood by looking at figure 5.5 (bottom), where the low spatial frequency results in an approximately uniform stimuli in the receptive field area. Also consistent with the

analytical results, less suppression was seen for higher spatial frequencies. Most variation in size tuning was seen for the frequency 0.04 cpd, while for higher frequencies, only a few of the measured cells show size suppression. In short, some neurons display surround suppression and some do not, depending on the combination of grating spatial frequency and filter size.

V1

The next question then becomes how this variation of LGN responses affects the size tuning of V1. To address this, the above LGN output was used to drive the V1. As for previous size tuning measurements, up to ten neurons from each population were analysed, depending on the number of cells within a 70 µm radius of the visual center. The response of a representative example neuron to all stimuli is shown in figure 9.19. The figure shows the response to increasing sizes of patch gratings. The

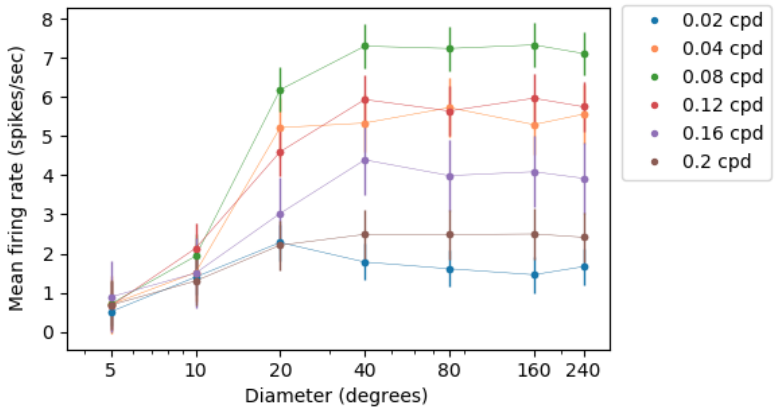

Figure 9.19: Size tuning of a representative V1 neuron (id = 109362, class = i23Htr3a) from the M1 Allen model. The visual stimuli was patches of drifting grating. The mean is calculated over 20 trials for 3-second simulations. Error bars show the standard deviation of the mean.

differences in scale of activity between the curves comes from the neuron's preference for different spatial frequencies. Only the curve for 0.02 cpd show a small amount of surround inhibition, while the remaining frequencies show no suppression. Results for all recorded neurons for all stimulus frequencies are included in the appendix, in figures B.28, B.29, B.30, B.31, B.32 and B.33. The general trend of the remaining neurons is similar to that in figure 9.19, with a small degree of size suppression seen for the frequency 0.02 cpd, but no suppression observed for higher frequencies.

In contrast, the experiments with patch gratings described in chapter 7, which all used different spatial frequencies, including the peak frequencies for each cell, all found cells with surround suppression. In conclusion, the surround effect which is observed experimentally for patch gratings in V1 is not reproduced by this particular M1 Allen model.

9.3 DoG Filters with Various Parameters

Up until now, only one set of parameters have been used for the DoG filter in equation (5.6) when running simulations. These parameters are $A_c = A_s = 1.73$ and $\sigma_s = 2.45\sigma_c$. For this case, no surround suppression is observed in V1 when using patch grating stimuli. To look into whether this is a result of the choice of DoG parameters, four new DoG filters were tested. Size tuning curves were then obtained for V1 cells in response to the same patch grating stimuli.

First, three filters with differently sized inhibitory Gaussians were tested. The scaling factor, C, in $\sigma_s = C\sigma_c$ was set to 5, 10 and 50 in the three cases. For the fourth DoG filter, the amplitudes were changed to $A_s = 0.5Ac$, while the size of the inhibitory Gaussian was the same as in the M1 model. As in the M1 model, all the amplitudes have been scaled so that peak LGN activity matches the peak activity in the original model (method 8.1.2). The new LGN spatial filters then have parameters

- $A_c = A_s = 1.2$ and $\sigma_s = 5\sigma_c$
- $A_c = A_s = 1.06$ and $\sigma_s = 10\sigma_c$
- $A_c = A_s = 1$ and $\sigma_s = 50\sigma_c$
- $A_c = 1.5$, $A_s = 0.5A_c$ and $\sigma_s = 2.45\sigma_c$.

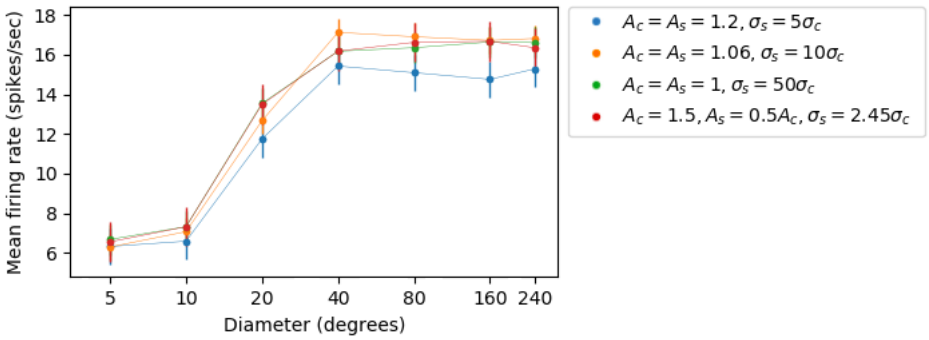

Figure 9.20: Size tuning curves for an example V1 neuron (id = 23428, class = e4). For each tuning curve, a different set of parameters for the DoG filters, defined by equation (5.6), is used. The parameters of each filter are indicated by the labels.

The size tuning of an example V1 neuron for the four new DoG filters is shown in figure 9.20. Similar curves for another four neurons are given in the appendix. Like for the M1 V1 model, all tuning curves show an increase in firing rate with diameter up to a point where the activity stabilises. This suggests that there is little or no surround effect in V1 for patch grating stimuli for any of the tested DoG filters.

9.4 Contributions to Size Tuning in V1

9.4.1 Effect of LGN input

So far, surround suppression has been observed in the M1 V1 for flashing spots, and in some cases for patch gratings with low spatial frequency (0.02 cpd). The suppression was generated after the spatial filters of the LGN, which transmit signals to V1, had been changed to describe center-surround receptive fields. If the receptive field of an Allen V1 cell is mainly inherited from the LGN cells transmitting to it, the suppression seen in V1 is arguably produced by an on-off antagonism in its receptive field. To illustrate this point, the sum of receptive fields of the M1 LGN cells relaying to each of two example M1 V1 cells is shown in figure 9.21.

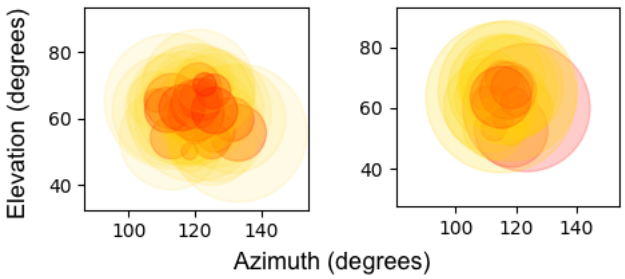

Figure 9.21: Illustration of the receptive fields of the M1 LGN neurons innervating two example V1 neurons. Left: LGN neurons relaying to V1 neuron 23428 (class = e4) consist of sON and tOFF cells. Right: LGN neurons relaying to V1 neuron 25515 (class = e5) consist of sOFF and tOFF cells. The two colours represent the center and surround fields of the M1 LGN neurons. On-regions are marked in yellow, while off-regions are marked in red. Azimuth and elevation refers to the horizontal and vertical angular coordinates in visual space.

Figure 6.2 in section 6.1 shows how the LGN cells transmitting to each V1 neuron are selected from two elliptical sub-regions in visual space. However, because of the LGN cells' large receptive fields, the sum of of the fields in the M1 model construct almost circular center-surround structures. As described in section 4.2, center-surround antagonisms can explain suppression in response to uniform spots, but fails to account for suppression in response to patch gratings for some spatial frequencies (see section 5.3.3). Furthermore, the suppression is most pronounced for low frequencies [95]. This description is consistent with the size tuning behaviour of M1 V1 cells, suggesting that the V1 receptive fields are in fact similar to the sum of M1 LGN receptive fields illustrated in figure 9.21. If V1 cells have mainly linear receptive fields with an on-off antagonism, this would explain the size tuning behaviour of the V1 cells, in that surround suppression is mostly seen for uniform spots rather than patch gratings.

However, the receptive fields of V1 cells may also be shaped by the recurrent connections between the V1 cells. To see how the spatial summation of V1 cells is influenced by network effects, size tuning was measured for a version of the M1 V1 model without recurrent connections.

9.4.2 Effect of recurrent connections in V1

In contrast to findings in the Allen model, experimental recordings in V1 show surround suppression in response to patch gratings for a range of different spatial frequen-

cies, including the cells' preferred frequencies. This implies that there are suppressive mechanisms in the real mouse's visual system not accounted for by the Allen model. What these mechanisms are, is, as seen in section 4.1.1, far from established. Some argue that suppression partly arises from within the V1 network itself, as a result of inhibition from other neurons [75, 140, 141].

To explore network effects on size tuning in V1, the M1 model was run without recurrent connections, using increasing sizes of patch gratings (see $M1_{\text{norec}}$ model in table 8.2). The same neurons as previously selected for size tuning were initially analysed. For the resultant plots, showing the activity of different cell types, see figure B.38 in the appendix. Tuning curves were then compared to those obtained from the M1 model with recurrent connections. For this, six neurons were selected from each class that receives LGN input. These are the inhibitory Htr3a class in LI, and the excitatory and Pvalb classes in LII/III to LVI [12]. Comparisons for all 54 neurons are shown in figure B.39 and B.40 in the appendix. Results for four example neurons, illustrating the different responses observed, are shown in figure 9.22.

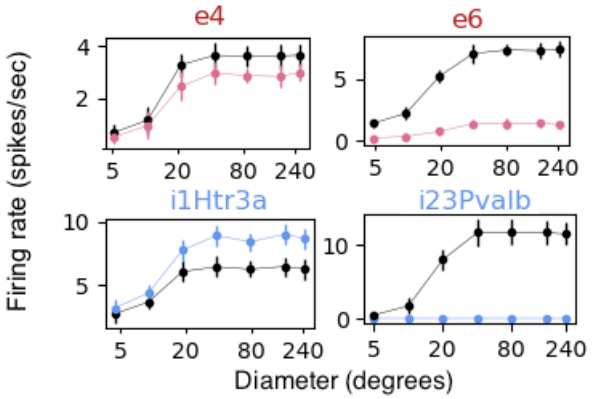

Figure 9.22: Comparison of size tuning for the M1 V1 model with (black) and without (red / blue) recurrent connections. The results for four representative cells are shown, illustrating the different responses displayed by the analysed cells. The tuning is measured in response to patch gratings with spatial frequency of 0.08 cpd.

The degree to which recurrent connections influence size tuning vary across layers and cell type. Without V1 connections, the activity depends on cell parameters such as the firing threshold, as well as the amount of LGN input, the latter of which LIV receives the most [39, 110]. Many V1 cells have similar tuning with and without recurrent connections, such as the e4 class (figure 9.22, top left). Other classes show only weak or no activity without V1 connections, such as the e6 and i23Pvalb cells, exemplified in figure 9.22. Variation was also seen within the different classes.

For cells which are active both with and without recurrent connections, the shape of the tuning curves are similar between the two cases. This could be explained by looking at the connectivity scheme in the model, which is described briefly in section 6.1. Whether connections are established depends, among other factors, on the distance between cells, where closer cells are more likely to be connected [12]. The probability distribution is given by equation (3.12), where the parameter σ_r takes different values between 85 and 120 µm, depending on cell types [12]. Using the mapping of visual to physical space in the Allen model [12], this corresponds to $6^\circ - 8.4^\circ$ azimuth and $3.4^\circ - 4.8^\circ$ elevation. Distance is only measured across the cortical plane, while the cortical depth was not taken into account [12]. As a result of using this

distribution, connections are established mainly between neurons closely positioned in the horizontal cortical plane. Therefore, only cells with similarly positioned receptive fields are connected. Furthermore, the tuning curves of all recorded cells reach maximum response at a similar patch size, of around 20° to 40°, suggesting similar receptive fields sizes. If connections are only between neurons with similarly sized and positioned receptive fields, this can explain why stimuli from regions far away from the visual center do not seem to influence the response of the neurons, neither by suppressing nor enhancing their activity.

In contrast, inhibitory surrounds up 92.5° in diameter have been observed in the real mouse V1 for patch gratings, while the median size producing maximum response was found to be 28.9° diameter [116]. Other studies report suppressive effects over similar ranges [64,115], suggesting that suppressive effects extend a great deal further than the intracortical connections in the Allen model.

The idea that connection lengths in the Allen V1 may be too short to account for far extending suppressive effects does not necessarily imply that the model should have longer intracortical connections. Most[1] connection lengths between the Allen V1 cells are based on experimental observations [12], and it has in fact been argued that connection lengths in real mice V1 are too short to account for at least the most remote surround effects [73]. On the other hand, recordings made by Adesnik et al. [115] suggest the existence of long reaching horizontal connections in LII/III, from the excitatory pyramidal cells to the inhibitory Sst neurons, which could effectively give the Sst neurons receptive fields spanning up to about 900 µm in diameter [12]. (See section 7.1.2 for a short description of this study.) In other words, the horizontal connections from excitatory neurons to Sst neurons may extend further than those of the Allen model. The same study also found that while the activity of Sst neurons increased with stimuli size, other neuron types displayed surround suppression [115]. Hence, increased inhibition from Sst neurons for large stimuli was suggested as a probable source of non-linear suppression in V1 [115].

In summary, recurrent connections in the Allen V1 do not seem to generate either surround suppression or enhancement, as connections are restricted to cells that have receptive fields of similar size and position. While some experiments support this connectivity scheme [12,73], others find evidence of long reaching horizontal connections in mice V1 [115] which may be able to account for some of the missing suppression in the Allen network.

9.5 Modelling Nonlinear Surround Suppression

Different models have been used to generate and describe surround suppression in the visual pathway. The normalisation model was first introduced as a descriptive model for nonlinear inhibitory phenomena observed in the cortex [142]. These phenomena included the saturation of neuron activity in response to increasing contrast, as well as nonspecific suppression evoked by a range of stimuli when superimposed on the preferred stimulus [142]. According to the normalisation model, the activity of a neuron is determined by the neuron's individual response, divided by the summed activity of a population of neurons [129].

[1]Not all connections in the V1 are based on experimental observations, due to lack of data. However, the connections most relevant for this section, from E to Sst cells, are based on experiments [12].

9.5.1 Divisive normalisation in LGN

Some have suggested that much of the nonlinear suppression seen in V1 first arises in the retina or LGN [63, 65, 79]. It was therefore interesting to see how nonlinear suppression in the LGN might affect the size tuning of V1 in the Allen model. The LGN was therefore modified from the original model to one which normalised the computed firing rates via equations 8.1 and 8.2. Specifically, for each time step, the firing rate of each cell was normalised based on the net LGN activity at that time. The LGN was then driven with patch gratings, for which it shows a high degree of size suppression, as seen in figure 9.23. Similar tuning is also seen in V1, as shown in figure 9.24. This points to divisive normalisation as a possible option for modelling nonlinear suppression in the Allen model.

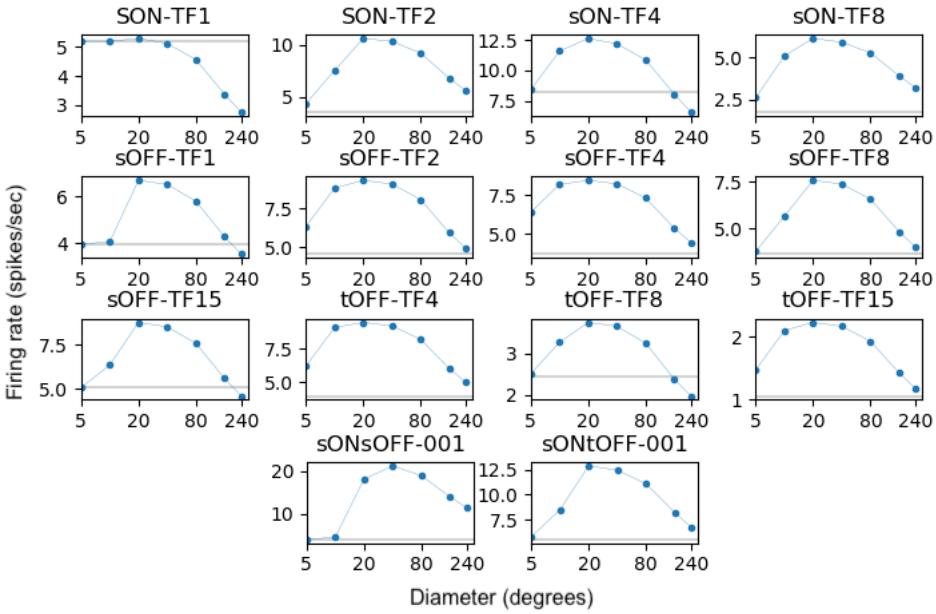

Figure 9.23: Size tuning (blue) of one LGN neuron per cell class from the normalised model. The spontaneous rate is marked by the grey line. The stimuli were patch gratings with spatial and temporal frequencies of 0.04 cpd and 4 Hz.

Note however that the normalisation approach taken here is a simple one, used primarily to demonstrate the suppressive effect and the fact that the property is transferred to V1. Compared to the strong suppressive effect in figure 9.23 and 9.24, experimental observations from mouse V1 suggest a weaker suppressive effect for patch grating stimuli (see chapter 7). This could for example be achieved by choosing a higher value of β in equations (8.1) and (8.2). Additionally, instead of summing over all LGN cells in the denominator of equation (8.1), a pool of cells could be chosen based on for example proximity [70] or similarity between cells, effectively creating a suppressive field for each neuron. This could be used to ensure that the response to large patch gratings do not suppress the activity below the spontaneous rate in the way demonstrated by the sON-TF1 cell in figure 9.23. For instance, a pool of cells could be created, consisting of only cells with similar tuning and spontaneous rates (e.g. other sON-TF1 cells).

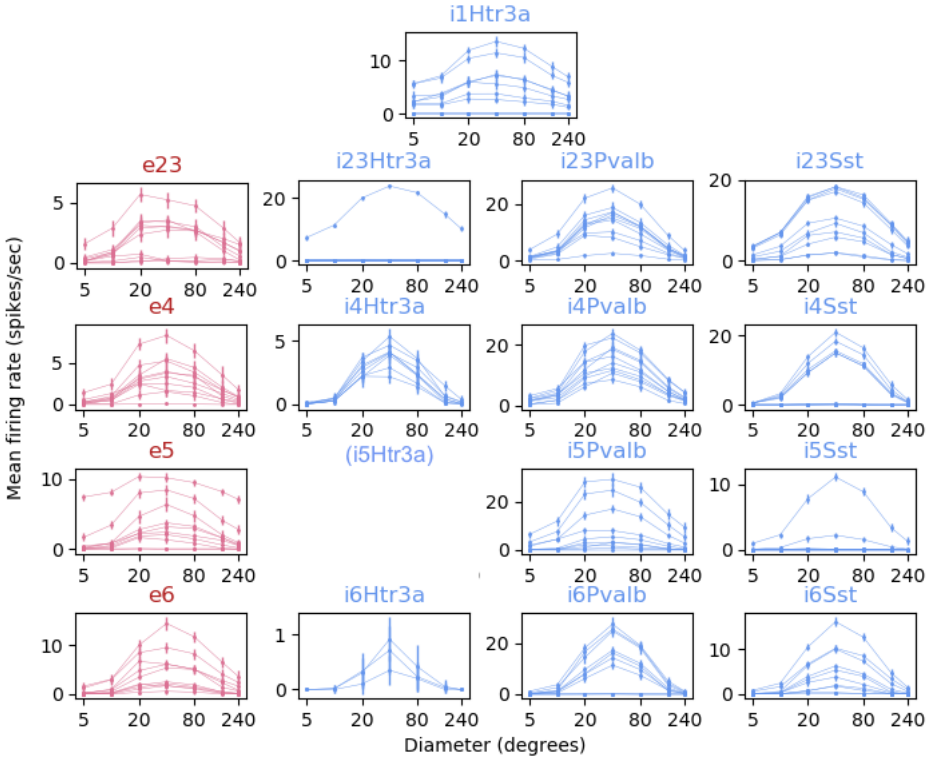

Figure 9.24: Size tuning of V1 neurons from the Allen model with normalised LGN responses. The plots show average measurements over 20 trials of 3-second simulations, with error bars representing the standard deviation. The visual stimuli were patch gratings with spatial and temporal frequencies of 0.04 cpd and 4 Hz. Up to ten neurons from each population are included, depending on the number of cells within 70 µm of stimuli center. For the i5Htr3a class, there were no cells within this radius. The different y-axes limits are chosen to show the tuning shape of all the populations. For an overview of how the degree of activity vary between populations, look to appendix figure B.41, for which the same scale is used for all subplots.

Alternatively, instead of applying the normalisation after computing the firing rates, one approach is to perform normalisation directly on the linear output of the receptive field, before applying the rectification [70]. Bonin et al. [70] used this to model suppression in the LGN. In their description, however, the sum in the denominator of equation (8.1) is substituted with a suppressive field, which depends on the standard deviation of contrast in the local region [70]. Hence a patch grating would evoke nonlinear suppression in the model. This method has been shown to successfully describe nonlinear properties of the cat LGN [70].

In addition to explaining nonlinear size suppression, divisive normalisation has also been shown to describe the saturation of responses to increasing grating contrasts, and nonlinear suppression caused by masking (superimposing another stimuli on the original one), both of which have been observed in the mammal LGN [70]. A number of mechanisms may underlie the processes which effectively normalise the LGN responses, including, as described in section 4.1.1, suppression inherited from the retina, inhibition from other LGN cells or feedback from V1. While findings that

70% of suppression in the cat LGN is accounted for by the retina [63] suggest the first, stimulation of mouse V1 cells have been found to suppress LGN cells [80], implying that both feedforward and feedback mechanisms are involved.

Limitations of the normalisation model

A critique of the normalisation description is that on its own it does not explain surround facilitation in response to weak center stimuli [125], a phenomenon reported by multiple studies [143, 144]. However, Cavanaugh et al. [73] argue that facilitation from surrounding regions are in fact just excitatory responses from areas within the classical receptive field with lower sensitivity to stimulation. In other words, stimulation of these areas alone only gives subthreshold responses [73]. It is therefore possible that the facilitatory effects are generated by a different mechanism than the suppressive effect of the surround. Cavanaugh et al. [73] further show how facilitation from peripheral regions can be effectively described using a contrast dependent receptive field size, showing that facilitatory effects could be incorporated in a variation of the normalisation model [73].

Considering the V1 in particular, another limitation relates to findings of orientation dependent suppression in real mice [64]. For example, when presenting a center patch grating surrounded by another grating, V1 suppression has been found to depend on the orientation of the latter [64] (see section 4.1.1). In particular, weaker suppression has been observed when the surround grating is oriented perpendicular to the center grating, than when oriented along the same axis [64]. In contrast, the suppression generated by normalising the Allen LGN responses is independent of orientation. In fact, a lack of orientation selectivity of LGN suppression in the cat [70] has been used to argue that LGN suppression is unable to fully account for suppression in V1, and that other areas must be involved. However, while the Allen LGN cells are independent of orientation, orientation-tuning has been shown to be present in the mouse LGN [52], although less so than in V1 [64]. One can therefore not exclude the possibility that orientation specific suppression in V1 partly stems from the LGN in mice. If so, orientation selectivity could be incorporated into LGN cells by for example modifying the DoG filter to an elliptical shape, oriented along some specified axis. In this case, the neurons contributing to suppression would have to be sampled based on orientation preferences, rather than using the sum of all LGN neurons like in equation (8.1).

Another property that is not described by the current model is delayed suppression. This delay has also been used to argue that other sources, besides the LGN, give rise to suppression in V1 [145]. Different delays have been recorded for surround suppression in V1 [145], including two separate delays found in the same study, indicating that more than one mechanism is involved [64].

In addition, it has to be emphasised that normalisation is a descriptive model which does not explicitly model the mechanisms behind nonlinear suppression. These mechanisms could for example include inhibitory connections within LGN or feedback from V1. To model these connections, one would have to replace the LNP filter with a mechanistic cell model such as the GLIF cell. Recurrent connections could for example be incorporated into an LGN network and adjusted to obtain suppression. Such an LGN description has in fact been developed by Mobarhan et al. [146]. In particular, the model is successfully able to generate surround suppression via network properties [146]. The study therefore gives an idea for how a mechanistic alternative to the Allen LGN could be implemented. That being said, a normalisation description

may be further used to represent nonlinear suppression in the Allen LGN, and provide insight into how the phenomena is transmitted to V1.

Chapter 10

Temporal Frequency Tuning of the Allen Model

In previous chapters, only the spatial tuning of the Allen model is explored. A second defining property of a visual stimulus is its temporal attributes. Similarly to other animals, mice are selective regarding the temporal features of stimuli, and multiple experimental studies have mapped the temporal frequency tuning of mice neurons [16, 116, 120–122, 147].

In section 5.1, a mathematical representation of a spatio-temporal receptive field is presented. The structure of the temporal filter component gives rise to the temporal tuning properties of the neuron. The Allen LGN cells, as described in section 6.2, consist of separable spatio-temporal filters, with a temporal kernel given by equation (5.5), as illustrated in figure 6.4. The kernel changes from positive to negative with time since stimulus, τ. This gives a preference for stimuli where the luminosity at a given location changes with time [50], for example drifting gratings. The optimal temporal frequency of the grating further depends on the shape of the filter.

The individual cell class parameters of the temporal filters in the Allen LGN [12] have been optimised on the basis of experimental cell recordings by Durand et al. [16] and Piscopo et al. [52]. However, the Allen documentation does not so far include tuning curves for all classes or the distributions of the preferred frequencies across the complete LGN and V1 models. In the following, the temporal tuning curves of the LGN and V1 cells will be presented. The tuning results are subsequently used to create distributions of preferred frequencies across populations of cells. The temporal selectivity of the Allen neurons are then compared to experimental findings.

The tuning properties were obtained using full-field drifting gratings of various temporal frequencies (see method 8.3 and 8.6) to drive the Allen model. Results for the M1 model is presented in this chapter. However, the original and M1 models were found to have the same temporal tuning (see appendix B.42 and B.45). This is because the LGN filters are spatio-temporally separable [12] (see section 6.2). Hence, a modification of these spatial filters from single Gaussians (original model) to DoG functions (M1 model) do not affect the temporal tuning.

10.1 LGN

The tuning curves in figure 10.1 show the firing rate dependence of LGN neurons on the temporal frequency of a grating. One neuron of each cell type is presented. For reference, the Allen model documentation includes a similar plot for the sON-TF8 class [12], displaying the same frequency tuning as found for the sON-TF8 in figure 10.1. Peak frequencies for each cell type are shown in figure 10.2.

Note that although the preferred frequencies in general agree with what is indicated by the cell class (e.g. TF1 denote a preference for 1 Hz), there is some inconsistency for the classes sON-TF1, sOFF-TF1, sOFF-TF2, sOFF-TF8 and tOFF-TF8. This could be due to an actual difference in the models and the classes they were based on [16], or it could be the result of different ways of measuring activity. Here, the mean activity over time, F0, is used, although some discrepancy was also found when using F1 as a measure (appendix B.43).

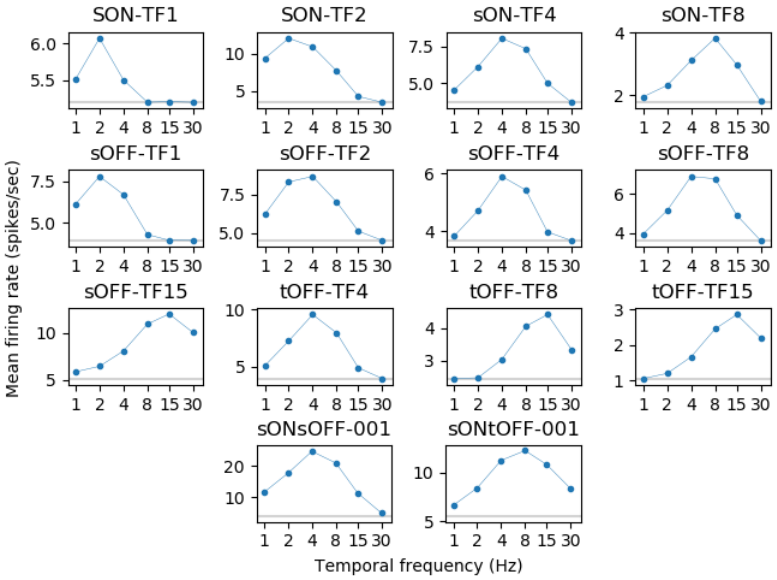

Figure 10.1: Temporal frequency tuning of LGN neurons. One example neuron from each cell class is included. The spontaneous rates are marked by the grey lines.

Figure 10.3 gives the total distribution of preferred frequencies across all LGN neurons. The median temporal frequency is 4 Hz, to which 48% (8385/17 400) of cells responded optimally. In comparison, Durand et al. [16], whose findings are used for fitting the LGN models, found that most cells responded optimally to 8 Hz in awake mice vs 4 Hz in in anaesthetised mice. The study reported mean peak frequencies of 6.27 Hz and 4.07 in the two states respectively. The values were obtained using gratings of 1, 2, 4, 8 and 15 Hz, comparable to the values used here (1, 2, 4, 8, 15 and 30 Hz).

Others have also looked at temporal tuning in the mouse LGN, reporting optimal frequencies of 3.8 Hz (mean) [113] and 3.2 Hz (median) [28]. However, both studies used anaesthetised animals, which were found by Durand et al. [16] to prefer significantly lower frequencies than the awake animal. If the Allen LGN is compared to the awake animal only, the optimal frequencies of the model are skewed towards slightly lower frequencies than in real mice.

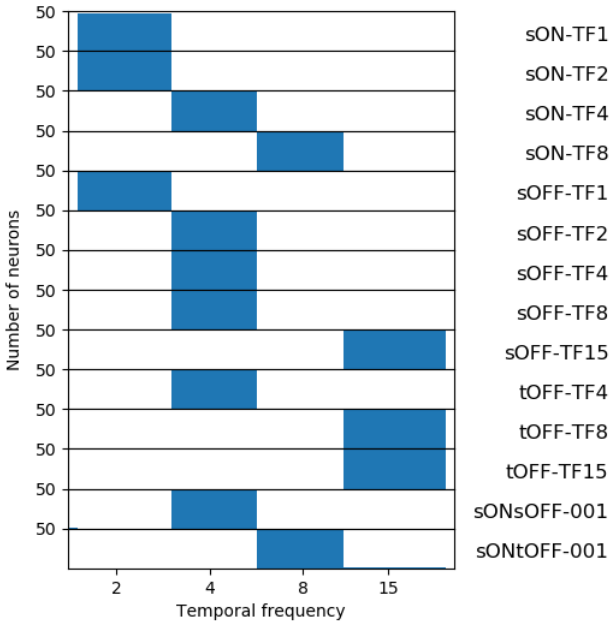

Figure 10.2: Optimal temporal frequencies for the different classes of LGN neurons, when driven by gratings of frequencies 1, 2, 4, 8, 15 and 30 Hz. 50 cells were analysed for each class. Not included in the figure are frequencies 1 and 30 Hz, to which none of the neurons showed a preference.

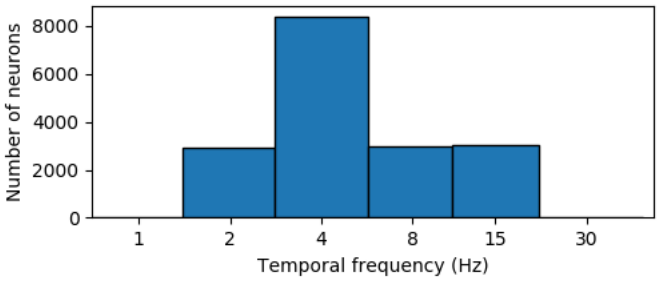

Figure 10.3: Preferred temporal frequencies for all neurons in the LGN.

10.2 V1

Temporal frequency tuning curves for 10 neurons of each cell type are presented in figure 10.4. The peak temporal frequencies of 50 neurons from each cell type is shown in figure 10.5. To represent the full V1 network, the preferred frequencies of 1000 randomly sampled neurons are shown in figure 10.6. In all cases, the unresponsive neurons have been excluded from the analysis. The recorded cells all display band-pass tuning, and nearly all V1 neurons prefer the temporal frequency of 4 Hz.

The study used for fitting Allen LGN parameters also includes recordings from V1 cells [16]. They found that the proportion of V1 cells responding optimally decreases with grating frequency, with mean preferred frequencies of 2.99 Hz (awake) vs 2.8 Hz (anaesthetised) [16]. A number of different preferred frequencies have also been reported by other studies on mice V1 [116, 120–123]. Both anaesthetised and awake mice have been used in these studies, though this was found by Durand et al. [16]

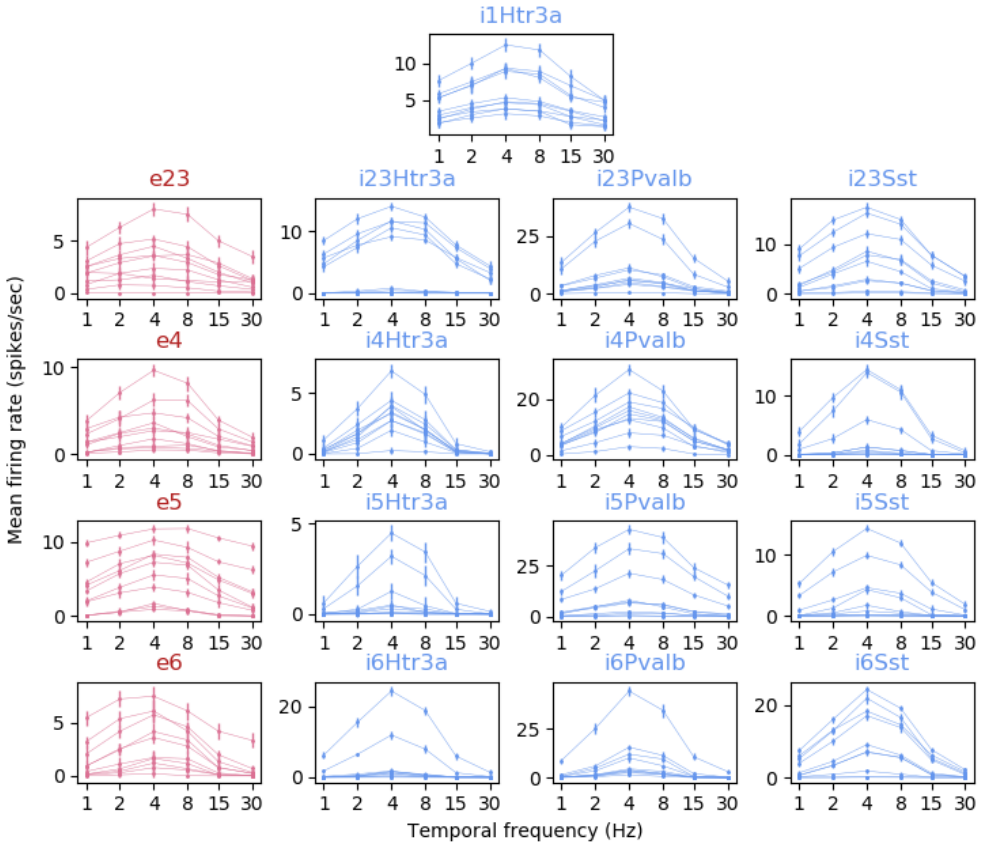

Figure 10.4: Temporal frequency tuning of V1 neurons in the M1 Allen model. Each subplot shows 10 responsive neurons from the respective classes. Curves show the mean over 20 trials of 3-second simulations, with error bars indicating the standard deviation. The different y-axes limits are chosen to show the tuning shape of all the populations. For an overview of how the degree of activity vary between populations, look to appendix figure B.47, where the same scale is used for all subplots.

not to influence temporal tuning in V1. In one experiement, Niell and Stryker [123] observed that neurons in LIV preferred significantly higher frequencies compared to other layers, with a mean peak of around 4 Hz vs around 2 Hz. The overall median peak frequency was found to be 1.68 Hz [123]. LeDue, Zou and Crowder [121] recorded optimal frequencies in the range 0.25 - 8 Hz, with a mean of 1.77 Hz. Gao, DeAngelis and Burkhalter [122] reported that about half the V1 neurons in LII/III - LVI had low-pass properties, with a median preferred frequency of 1.2 Hz, while the remaining neurons showed band-pass tuning, with a significantly higher median at 1.9 Hz. Van den Bergh et al. [116] and Andermann et al. [120] measured median optimal frequencies of 3.4 Hz (all layers) and 3 Hz (LII/III) respectively. These reported peak frequencies in V1 are, on average, lower than those reported from LGN [16, 28, 113]. Correspondingly, Durand et al. [16] found that V1 preferred significantly lower frequencies than LGN.

When comparing these recordings with the Allen V1, there is a shift towards lower preferred frequencies in real mice, going from LGN to V1, which is not displayed by the Allen model. In contrast, the Allen V1 cells have the same median peak frequency

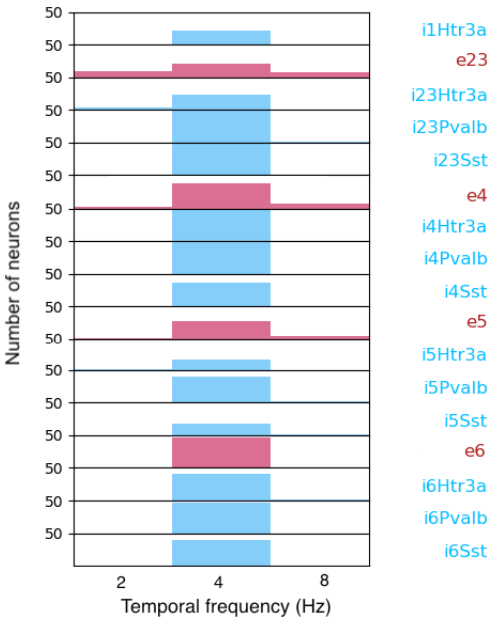

Figure 10.5: Optimal temporal frequencies of V1 neurons in the M1 Allen model. The labels represents the different V1 classes. For each class, 50 cells are analysed, and the preferred frequencies of the responsive neurons are shown. Gratings of temporal frequencies 1, 2, 4, 8, 15 and 30 Hz are used as visual stimuli. The frequencies 1, 15 and 30 Hz, to which none of the selected neurons showed a preference, are omitted from the figure. Preferred frequencies are based of mean firing rates over 20 trials of 3-second simulations.

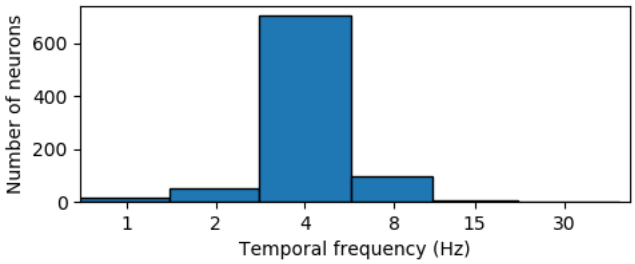

Figure 10.6: Preferred temporal frequencies for 877 responsive cells out of 1000 randomly drawn V1 neurons from the M1 Allen model.

of 4 Hz as the Allen LGN. This could imply that the temporal tuning in the Allen V1 is largely determined by the temporal tuning of the Allen LGN. Similar results were found for spatial frequency tuning in chapter 9. As a result, the optimal temporal frequencies in the Allen V1 are higher than those observed in the real mouse V1.

The experimental studies also show that different V1 neurons display a range of preferred frequencies. This variety is not exhibited by the V1 neurons in the Allen model, which are relatively homogeneous (figure 10.4). The distribution of preferred frequencies in the V1 model is substantially narrower than in the LGN model, with 80 % (704/877 responsive cells) preferring the 4 Hz grating, compared to 48 % (8385/17 400 responsive cells) in the LGN. This could be due to preferences inherited from

the LGN cells averaging out, as several different LGN neurons transmit to each V1 cell. There is some dependence on temporal tuning preferences in the connectivity between LGN and V1 in the Allen model, where some LGN classes are more likely to be connected to certain V1 cell types [12]. However, LGN cells with different temporal preferences still project the same V1 cells. As an example, one single neuron (id = 25515) receives input from LGN cells preferring the frequencies 1, 2, 4, 8, and 15 Hz [148]. Recurrent connections likely also influence the temporal selectivity in V1, an effect which could be looked into by for example turning off recurrent connections in V1 and repeating the temporal tuning analysis.

In summary, Allen V1 cells prefer higher temporal frequencies than in real mice. While the LGN and V1 models have the same optimal frequency, real mice have a shift towards lower preferred frequencies from LGN to V1. Furthermore, mice V1 display a wider variety of optimal frequencies that the V1 model. The frequency tuning of the Allen model can largely be explained by LGN preferences being transmitted to, and averaging out in V1. Therefore, there may be network effects in the real mice influencing the temporal tuning which are not accounted for by the model.

Chapter 11

Summary, Conclusion and Future Work

Spatial and temporal properties of the Allen model have been explored. Specifically, spatial and temporal frequency tuning and size tuning of LGN and V1 cells have been measured. For this, three types of visual stimuli have been used: full-field gratings, patch gratings and flashing spots. The results have then been compared to findings from experiments on real mice.

11.1 Size and Spatial Frequency Tuning

First, the behaviour of the original Allen model is considered. This is followed by a discussion of the modified versions. The results from both models are compared to experimental recordings. To make sure that modifications were based on consistent experimental findings, a selection of relevant studies on real mice have been presented in chapter 7.

11.1.1 Original model

LGN

The Allen model consists of two modules: the LGN and V1. The first step for running simulations is to present visual input to the LGN. The resulting output is then used to drive the V1 model. We start by considering the LGN responses. Since the LGN cells are modelled as Linear-Nonlinear-Poisson filters, their behaviour is largely that of a linear receptive field. In section 5.3, the spatial tuning of such a field is either derived or described for the relevant stimuli. In chapter 9, when looking at simulation output from the Allen LGN, none of the analysed LGN cells showed size suppression in response to flashing spots or patch grating stimuli. This is in contrast to previous studies with flashing spots, which reported size suppression in a large proportion of cells [52, 83, 113]. Furthermore, all analysed Allen LGN cells showed low-pass spatial frequency tuning, inconsistent with the band-pass tuning reported by previous experiments [16, 28, 52, 83, 113].

Both the low-pass tuning and lack of size suppression in the Allen LGN can be explained by the fact that its receptive field consists of a simple 2D Gaussian function. Since there is no on-off antagonism, there is no mechanism to cause preference for stimuli which are non-uniform in space, as explained in section 4.2. This agrees

with the mathematical predictions of 5.3. Experimental studies, on the other hand, found that LGN cells had receptive fields with on-off antagonisms in center-surround structures [52, 83, 113].

V1

In contrast to the LGN, the Allen V1 comprises a network of GLIF neurons. Each neuron may receive input from both LGN cells and other V1 cells. In addition, each cell is subject to a constant background activity. The behaviour of V1 is therefore more complex than that of LGN and not easily predicted. Nevertheless, similar spatial frequency and size tuning was found in V1 and LGN. The analysed V1 cells showed no surround suppression for either patch gratings or flashing spots. In contrast, studies report that most cells showed size suppression for patch-gratings [48, 64, 82, 115, 116]. One study using flashing spots also found size suppression for a large proportion of cells [117]. Looking at spatial frequency tuning, most Allen V1 cells showed low-pass behaviour, with peak response at the lowest presented spatial frequency (0.005 cpd). Some V1 cells also preferred slightly higher frequencies, something which was not seen in LGN. Since most inhibitory cells also showed low-pass tuning, inhibition will be strongest at the lowest frequency. In some cases, this may cause another cell's activity to peak at higher frequencies, and could be the reason for the observed variation in preferred frequencies. In comparison to the Allen V1, studies on real mice found that the average cell showed spatial frequency tuning with band-pass behaviour [16, 116, 118]. To summarise the results, the original Allen model show low-pass tuning and fail to exhibit size suppression for both flashing spots and patch gratings. There is therefore a discrepancy between measurements in the model and in real mice, the latter which have been found to exhibit size suppression and band-pass tuning.

11.1.2 M1 models

Motivation

As described above, the size and spatial frequency tuning of the original Allen model is not consistent with that found in real mice. The further aim was therefore to make the model more realistic in terms of these properties. Hence, modifications for reproducing the phenomena in the model were explored. Since the differences observed between the Allen V1 and real mice could potentially be due to the LGN input, the first step was to look at the LGN. Accordingly, the LGN receptive fields, consisting of a center-only structure, were replaced by the center-surround structure found in real mice. The center-surround fields were represented by a DoG function, with the size of the surround field chosen to give a median preferred spatial frequency of 0.08 cpd, equal to that measured by Durand et al. [16]. The center and surround amplitudes were chosen to give maximum suppression, although realistically, the surround amplitude may be about half that of the center, as the mean suppression index in LGN has been measured to be 0.53 [83]. With this modification, the analytical linear prediction of the cell responses, as presented in section 5.3, follow experimental recordings by exhibiting size suppression for uniform stimuli and band-pass spatial frequency tuning. The next step was then to explore the effect of this modification on LGN and V1.

LGN

The properties of the M1 model were measured by repeating the same simulations as in the original model, both for LGN and V1. The M1 LGN cells showed band-pass spatial frequency tuning, with preferred frequencies in the range 0.041 to 0.201 cpd. Although the range is smaller, this is similar to the variation seen in real mice [16]. To create a wider range of preferred frequencies, the inhibitory surrounds of the DoG filters can be chosen so that there is a larger variation in size. For the size tuning measurements, the cells showed surround suppression for isoluminous flashing spots, with a preferred size ($\sim 10^{\circ}$) similar to that measured in real mice [52, 83]. For patch grating stimuli, the cells showed mixed responses, depending on the stimulus spatial frequency and the cell's DoG parameters. Some neurons showed varying degree of surround suppression, while some showed no surround effect. The LGN responses were consistent with the analytical predictions of a DoG receptive field to the various stimuli (full-field gratings, uniform spots and patch gratings).

V1

The M1 V1, like the M1 LGN, showed size suppression for flashing spots and band-pass spatial tuning, with a median preferred frequency of 0.08 cpd. As in the original V1, there was some variation in preferred frequency, although the majority of cells preferred the grating with frequency of 0.08 cpd. In comparison, V1 cells in real mice show a large variation in optimal spatial frequencies [16,118]. For patch grating stimuli, weak surround inhibition was found for spatial frequencies of 0.02 cpd, while no suppression was seen for frequencies 0.04, 0.08, 0,12, 0.16 or 0.2 cpd. In contrast, experimental findings of surround inhibition in V1 were found for gratings of various spatial frequencies, including the preferred frequency of single cells [64, 82, 115, 116], which for the Allen M1 V1 was around 0.08 cpd. To make sure the lack of suppression was not due to the choice of DoG parameters, four DoG filters with different parameters where tested, of which none showed surround suppression. To summarise, the M1 model was able to reproduce band-pass tuning and surround suppression for flashing spots. However, the model does not account for suppression in response to patch gratings.

11.1.3 Origin of tuning properties in V1

Analytical solutions go a long way in explaining the observed properties of the Allen model. As a filter model, the LGN largely follows the linear prediction of the spatial receptive field function (section 5.2). Also in the V1 network, the tuning properties are suggestive of a predominantly linear spatial summation. First, the cells show either low-pass tuning and no size suppression (original model) or band-pass tuning and size suppression for flashing spots (M1 model). This relationship between size and frequency tuning is as described by linear spatial frequency analysis in section 4.2. Second, the M1 model show weak suppression for low spatial frequencies, and no clear suppression for higher frequencies, which is similar to the behaviour of a DoG filter [95] (see section 5.3.3). In fact, when looking at the LGN neurons relaying to a single V1 cell, their receptive fields construct an almost center-surround structure.

For a single V1 neuron, there are two potential sources for influencing its tuning properties: LGN input and recurrent connections.[1] While the results above makes

[1] Each V1 cell also receives a background input, but since the strength of this input is a constant for each cell, this influence is not considered.

it seem as if the receptive field of a V1 neuron is purely inherited from the LGN neurons transmitting to it, many V1 neurons do not receive any LGN input, but still respond to stimuli. Hence, the recurrent connections in V1 must necessarily also play a role in shaping the cells' receptive fields. To explore this effect, additional size tuning simulations were run without recurrent connections in V1. The recurrent connections in the Allen V1 did not seem to qualitatively influence the tuning curve, except for giving rise to activity in otherwise silent neurons. This can be ascribed to the connections primarily being between cells at similar horizontal positions (position in the cortical plane). The connected V1 cells therefore have similar receptive fields, causing the effect of the recurrent connections to primarily be a scaling of the size tuning curve, rather than an influence on its shape. To summarise, the qualitative tuning properties of V1 seem to be largely determined by the LGN input, while recurrent connections have a limited role. As a consequence, the V1 tuning results are similar to that of the LGN cells.

Although only the size tuning of V1 cells was measured without recurrent connections, it could be interesting to do the same for spatial frequency tuning. V1 and LGN has the same mean optimal frequencies, however a lot less variation is seen across cells in V1. Tuning measurements without recurrent connections could therefore be done to check whether preferences average out due to network effects or because multiple LGN cells project to each V1 cell. From here, one could look into ways of generating more variation of preferred frequencies in V1. One suggestion is to modify the connections between LGN and V1, and consider the preferred spatial frequencies when determining which LGN neurons should transmit to each V1 cell. Connections could then be established mainly between cells of similar frequency tuning. Recurrent connections in V1 could also be adjusted according to the cells' preferred spatial frequencies.

11.1.4 Modelling nonlinear suppression

By exploring the M1 model, we have seen that center-surround receptive fields are able to account for band-pass tuning and size suppression for flashing spots in LGN, and that these properties are propagated further to V1. However, neither the original nor M1 models describe the size suppression seen in V1 for patch gratings. A likely explanation is that spatial summation in the Allen V1 seems predominantly linear, which cannot fully account for the suppression seen in response to patch gratings (see section 4.1.1). This type of suppression is therefore often referred to as nonlinear suppression.

The understanding of nonlinear suppression and how it originates in real animals is currently lacking [65], although recurrent connections and feedback inhibition between and within the LGN and V1 are thought to play a role [74, 76–79, 81, 149] (see section 4.1.1). It has also been suggested that the mechanism first arises in the retinal ganglion cells and is inherited by the LGN and V1 [63]. Nevertheless, the simulations from the original and M1 Allen model indicate that there exists some nonlinear mechanisms in the real mouse visual pathway which are not accounted for by the model. In the following, different approaches for modelling nonlinear suppression are discussed. This thesis explores one such method, divisive normalisation in LGN, as a way of generating suppression for patch gratings. Other options involve modelling recurrent connections in V1 or various feedback contributions, and will be suggested as potential topics for further research.

Divisive normalisation in LGN

Although no experimental studies on extraclassical surround inhibition were found for mouse LGN, the property has been found in the LGN of other animals, such as cats and monkeys [63, 65, 66, 79]. It was therefore interesting to explore how this can be reproduced in the Allen LGN and in what way it affects size tuning in V1. The LGN was modified to use a simple form of divisive normalisation, where the cell activity at each time step was scaled according to the net sum of LGN activity at that time, as given by equation (8.1). It was then shown that the resulting suppression is propagated from LGN to V1.

This suggests LGN normalisation as an alternative for modelling nonlinear suppression, although the form of normalisation would probably have to be adjusted to make it more realistic. One option is to adopt a scheme similar to that presented by Bonin et al. [70], where the response of a center-surround receptive field is divided by that of a larger, suppressive field, before a rectification is applied to give the final firing rates. This scheme has been shown to predict not only size suppression, but other nonlinear effects such as response saturation with increasing contrast [70].

One limitation of modelling nonlinear suppression in the Allen V1 as completely inherited from the LGN is the lack of orientation selectivity of the Allen LGN cells. Suppression of grating stimuli in the mouse V1 is sensitive to orientation [64], and although mouse LGN cells have been found to display some orientation selectivity [52], the lack of similar selectivity in other mammal LGNs has been used to argue that nonlinear suppression in LGN and V1 either arises partly independently [70] or at least undergoes some sort of modification from the LGN to V1 [63]. Nevertheless, a normalisation approach could be used to incorporate nonlinear effects in the Allen LGN, and its effect on size tuning and possibly contrast saturation are interesting topics to explore further.

While nonlinear suppression in LGN has been investigated for mammals such as the cat [66, 70] and monkey [63], similar experiments in mice are currently lacking. To be able to accurately compare the results of the Allen LGN with real mice, further experimental studies are required. It could therefore be useful to look at size tuning in response to for example patch gratings in the real mouse LGN, and compare the results with the findings of this thesis.

Inhibition via recurrent connections in V1

Another option for introducing suppression in the Allen model is through recurrent connections in V1. Rather than just describing observed results, this allows for modelling underlying physical mechanisms for suppression. Furthermore, properties which are not here accounted for by normalising the LGN responses, such as delayed or orientation dependent suppression, might be explained by V1 network effects. In real mice, opinions about the role of horizontal intracortical connections in surround suppression are divided. While some studies suggest connections are too short to account for size suppression [73], others have found evidence for long reaching horizontal connections from excitatory cells to the inhibitory Sst class in LII/III [115]. The increased activity in Sst neurons with stimuli size, together with findings of suppression in other cells types, suggest that inhibition from Sst neurons with increased stimuli size could play a role in nonlinear suppression [115]. In relation to the Allen model, horizontal connections from the E to Sst class could be extended to the same scale as found by Adesnik et al. [115]. In addition, the synaptic strengths could be adjusted so that, according to the findings by Adesnik et al. [115], the main driving force of Sst cells would be E cells from the same layer. By making these modifications, one

could explore if nonlinear surround suppression in the Allen V1 may be modelled in a mechanistic way, using properties observed in the real mouse V1.

Feedback and feedforward contributions

Some studies suggest that suppression arises through interactions between different areas in the visual pathway. For example, recordings have shown that stimulation of V1 cells has the potential to suppress responses in LGN [80]. Cortical inactivation has also been found to reduce LGN suppression [76–78]. Analogously, some studies point to feedback from higher cortical areas as likely sources of V1 suppression [145]. Since the Allen model consist of separate LGN and V1 modules, potential feedback effects were not explored in this thesis.

In the future, it could be interesting to replace the LGN filters with mechanistic neuron models, and extend the Allen network to include an interaction between LGN and V1. Similarly, Mobarhan et al. [146] has developed a network model for exploring the effects of cortical feedback on LGN. The network successfully generated surround suppression for both circular spots and patch gratings [146]. In future studies, the model by Mobarhan et al. [146] could be used as a starting point for determining the form of feedback connections in a modified Allen model. This feedback could include connections from V1 to LGN, or input from cells representing higher cortical areas.

11.2 Temporal Frequency Tuning

In addition to looking at spatial tuning properties of the Allen model, the temporal frequency tuning of the model was briefly explored. In the Allen model, LGN and V1 have the same average peak temporal frequency of 4 Hz. When comparing with experiments, there is a shift in preferred frequencies, going from LGN to V1 in real mice, not observed in the Allen model [16]. The preferred frequencies of Allen V1 cells were therefore higher than that observed experimentally [116, 120–123]. A suggestion for further research is to explore mechanisms which could cause this shift. Niell and Stryker [123] reported a significantly higher frequency preference in LIV compared to other layers [123]. Recall that LIV is the main target of LGN input [36]. This suggests the involvement of network effects within V1 or feedback from higher cortical areas in the shift towards lower frequencies.

Furthermore, experiments show large variability in temporal selectivity across V1 cells [16, 121–123]. This is in contrast to the Allen V1 network, which is relatively homogenous, with the vast majority of V1 cells preferring the same frequency of 4 Hz. This similarity across cells was found for both spatial and temporal tuning, suggesting less variety between cells in the Allen V1 network than in the real mouse. The LGN model on the other hand, has a much broader range of preferred frequencies. In the Allen model, established connections from LGN to V1 depend, to some degree, on temporal tuning properties [12]. Specifically, some LGN classes are more likely to be connected to certain V1 cells [12]. Still, multiple LGN cells preferring different temporal frequencies project to the same V1 cells. This can lead to preferences of LGN cells averaging out in V1. A connection scheme between LGN and V1 with a higher selectivity of temporal tuning may give rise to more variation in V1. Recurrent connections in V1 may also be adjusted to depend on temporal frequency preferences.

11.3 Conclusion

In this project, selected properties of the Allen model have been explored. The main focus has been on spatial tuning, namely spatial frequency tuning and size tuning. In addition, the temporal frequency tuning of cells has been analysed. The temporal tuning of LGN cells, developed and optimised according to experiments [16, 52], was first mapped. The same tuning analysis was then done for V1, where optimal frequencies were found to be somewhat higher than experimental values [16,116,120–123]. In addition, less variability was found in the Allen V1 than in experiments [16,121–123], something which also applied to the spatial tuning measurements.

Looking at the original Allen model, the spatial tuning results reveal inconsistencies with experimental evidence on mice, namely lack of size suppression and primarily low-pass tuning in both LGN and V1. Therefore, ways of improving the model were considered. In order to make sure that the modifications were grounded in literature, multiple experiments on real mice were reviewed. This resulted in two different approaches to modifying the model.

The first addressed the spatial structure of LGN receptive fields, which have been reported to have a center-surround structure [52, 83, 113]. The receptive fields of the Allen LGN model, in contrast, had center-only structures. In section 5.3, we see how a center-surround structure can account for properties such as band-pass tuning and size suppression of uniform stimuli. Indeed, when modifying the LGN spatial fields from Gaussian to DoG filters, these properties were generated in both LGN and V1.

For patch grating stimuli however, the first modification barely gave rise to size suppression in V1, in contrast to what has been observed in mice [48,64,82,115,116]. For this, some form of nonlinear suppression had to be introduced into the model. Hence, the LGN was modified to apply a simple divisive normalisation on the firing rates. The resultant suppression generated in the LGN model was shown to propagate to V1. Therefore, normalisation has potential as a descriptive model for nonlinear suppression in the Allen LGN. However, properties of suppression in V1, such as orientation selectivity [64] and different delays in onset [64,145], suggest the existence of multiple mechanisms contributing to the suppression in V1.

Another potential source of nonlinear suppression is recurrent connections within V1. In the Allen model, these connections are generally between cells with similar receptive fields, and do not seem to give rise to either surround enhancement or suppression. This is in contrast to one study, reporting long-reaching horizontal connections from excitatory pyramidal cells to Sst neurons in LII/III, with the former being the main source of input for the latter [115]. Adjusting the Allen V1 to integrate these findings could therefore be an option for further exploration of nonlinear suppression in Allen V1.

11.4 Further Research

Being a model based on real mice, the Allen model opens up opportunities to explore several properties and test the results against experiments. As the unanswered questions in this thesis mainly concerns nonlinear suppression, a few suggestions for exploring this property are made below. Suggestions for looking into spatial and temporal tuning behaviour are also proposed.

- Based on findings by Adesnik et al. [115], one could extend horizontal connections from E cells to Sst cells in V1, and scale the synaptic strengths so that horizontal connections from E cells becomes the main input to Sst cells. This can

further be used to test if large Sst receptive fields could introduce suppression in other V1 neurons.

- Further develop the normalisation model for nonlinear suppression in LGN, for example by integrating a suppressive field in combination with a center-surround structure, according to the scheme proposed in [70].
- Measure size tuning in real mice LGN in response to patch gratings of multiple spatial frequencies. Findings can then be compared to the behaviour of the Allen LGN, so that a more accurate model of possible nonlinear LGN suppression can be developed, for example in the form of divisive normalisation.
- Model contributions of feedback on surround suppression in LGN and V1. One could look into ways of modelling feedback from V1 or from higher cortical areas. A network model of LGN, developed by Mobarhan et al. [146], has been used to model the effect of cortical feedback on surround suppression. This description could therefore be used as a starting point for determining how to best model feedback effects in a modified version of the Allen model.
- Look for ways to increase the variation in V1 cells regarding spatial and temporal frequency tuning. The model could first be driven without recurrent connections for determining to what extent LGN input and network effects influence the frequency tuning. Connections within V1 or between LGN and V1 could then be adjusted according to the cells preferred frequencies.
- Look into sources of the experimentally observed shift towards lower preferred frequencies when going from LGN to V1 in real mice. The next step would be to explore how the Allen model could be modified to reproduce this behaviour.

Bibliography

[1] C. Teeter, R. Iyer, V. Menon, N. Gouwens, D. Feng, J. Berg, A. Szafer, N. Cain, H. Zeng, M. Hawrylycz, *et al.*, "Generalized leaky integrate-and-fire models classify multiple neuron types," *Nature communications*, vol. 9, no. 1, pp. 1–15, 2018.

[2] M. Arbib, ed., *The handbook of brain theory and neural networks*. MIT press, 2003.

[3] M. Nelson and J. Rinzel, "The hodgkin—huxley model," in *The book of genesis*, pp. 29–49, Springer, 1998.

[4] A. L. Hodgkin and A. F. Huxley, "The components of membrane conductance in the giant axon of loligo," *The Journal of physiology*, vol. 116, no. 4, pp. 473–496, 1952.

[5] A. L. Hodgkin and A. F. Huxley, "Currents carried by sodium and potassium ions through the membrane of the giant axon of loligo," *The Journal of physiology*, vol. 116, no. 4, pp. 449–472, 1952.

[6] A. L. Hodgkin and A. F. Huxley, "The dual effect of membrane potential on sodium conductance in the giant axon of loligo," *The Journal of physiology*, vol. 116, no. 4, pp. 497–506, 1952.

[7] A. L. Hodgkin, A. F. Huxley, and B. Katz, "Measurement of current-voltage relations in the membrane of the giant axon of loligo," *The Journal of physiology*, vol. 116, no. 4, pp. 424–448, 1952.

[8] A. L. Hodgkin and A. F. Huxley, "A quantitative description of membrane current and its application to conduction and excitation in nerve," *The Journal of physiology*, vol. 117, no. 4, pp. 500–544, 1952.

[9] W. Gerstner, W. M. Kistler, R. Naud, and L. Paninski, *Neuronal dynamics: From single neurons to networks and models of cognition*. Cambridge University Press, 2014.

[10] C. Rössert, C. Pozzorini, G. Chindemi, A. P. Davison, C. Eroe, J. King, T. H. Newton, M. Nolte, S. Ramaswamy, M. W. Reimann, *et al.*, "Automated point-neuron simplification of data-driven microcircuit models," *arXiv preprint arXiv:1604.00087*, 2016.

[11] A. Arkhipov, N. W. Gouwens, Y. N. Billeh, S. Gratiy, R. Iyer, Z. Wei, Z. Xu, R. Abbasi-Asl, J. Berg, M. Buice, *et al.*, "Visual physiology of the layer 4 cortical circuit in silico," *PLoS computational biology*, vol. 14, no. 11, p. e1006535, 2018.

[12] Y. N. Billeh, B. Cai, S. L. Gratiy, K. Dai, R. Iyer, N. W. Gouwens, R. Abbasi-Asl, X. Jia, J. H. Siegle, S. R. Olsen, *et al.*, "Systematic integration of structural and functional data into multi-scale models of mouse primary visual cortex," *bioRxiv*, p. 662189, 2019. doi: `https://doi.org/10.1101/662189`.

[13] D. Sterratt, B. Graham, A. Gillies, and D. Willshaw, *Principles of computational modelling in neuroscience*. Cambridge University Press, 2011.

[14] L. Benuskova and N. K. Kasabov, *Computational neurogenetic modeling*. Springer Science & Business Media, 2010.

[15] M.-P. Schallmo, A. M. Kale, and S. O. Murray, "The time course of different surround suppression mechanisms," *Journal of vision*, vol. 19, no. 4, pp. 12–12, 2019.

[16] S. Durand, R. Iyer, K. Mizuseki, S. de Vries, S. Mihalas, and R. C. Reid, "A comparison of visual response properties in the lateral geniculate nucleus and primary visual cortex of awake and anesthetized mice," *Journal of Neuroscience*, vol. 36, no. 48, pp. 12144–12156, 2016.

[17] M. A. Bedell, N. A. Jenkins, and N. G. Copeland, "Mouse models of human disease. part i: techniques and resources for genetic analysis in mice.," *Genes & development*, vol. 11, no. 1, pp. 1–10, 1997.

[18] N. Rosenthal and S. Brown, "The mouse ascending: perspectives for human-disease models," *Nature cell biology*, vol. 9, no. 9, p. 993, 2007.

[19] M. Baker, "Neuroscience: through the eyes of a mouse," *Nature News*, vol. 502, no. 7470, p. 156, 2013.

[20] T. A. Seabrook, T. J. Burbridge, M. C. Crair, and A. D. Huberman, "Architecture, function, and assembly of the mouse visual system," *Annual review of neuroscience*, vol. 40, pp. 499–538, 2017. doi: `https://doi.org/10.1146/annurev-neuro-071714-033842`.

[21] D. Simmons, "The use of animal models in studying genetic disease: transgenesis and induced mutation," *Nature education*, vol. 1, no. 1, p. 70, 2008.

[22] L. A. Remington and L. Remington, "Visual pathway," in *Clinical Anatomy and Physiology of the Visual System*, pp. 233–252, USA: Elsevier Butterworth-Heinemann, 2012.

[23] A. Monavarfeshani, U. Sabbagh, and M. A. Fox, "Not a one-trick pony: diverse connectivity and functions of the rodent lateral geniculate complex," *Visual neuroscience*, vol. 34, 2017.

[24] C. Watson, "Visual system," in *The Mouse Nervous System*, pp. 646–652, Elsevier, 2012.

[25] E. R. Kandel, J. H. Schwartz, T. M. Jessell, D. of Biochemistry, M. B. T. Jessell, S. Siegelbaum, and A. Hudspeth, *Principles of neural science*, vol. 4. McGraw-hill New York, 2000.

[26] J. Stirling, *Cortical functions*. Routledge, 2002.

[27] C. Galindo-Romero, M. Aviles-Trigueros, M. Jimenez-Lopez, F. Valiente-Soriano, M. Salinas-Navarro, F. Nadal-Nicolas, M. Villegas-Perez, M. Vidal-Sanz, and M. Agudo-Barriuso, "Axotomy-induced retinal ganglion cell death in adult mice: quantitative and topographic time course analyses," *Experimental eye research*, vol. 92, no. 5, pp. 377–387, 2011.

[28] J. Tang, S. C. A. Jimenez, S. Chakraborty, and S. R. Schultz, "Visual receptive field properties of neurons in the mouse lateral geniculate nucleus," *PloS one*, vol. 11, no. 1, p. e0146017, 2016.

[29] J. K. Mai and F. Forutan, "Thalamus," in *The human nervous system*, pp. 618–677, Elsevier, 2012.

[30] B. P. Covington and Y. Al Khalili, "Neuroanatomy, nucleus lateral geniculate," 2019.

[31] W. M. Usrey and H. J. Alitto, "Visual functions of the thalamus," *Annual review of vision science*, vol. 1, pp. 351–371, 2015.

[32] A. Pienaar, L. Walmsley, E. Hayter, M. Howarth, and T. M. Brown, "Commissural communication allows mouse intergeniculate leaflet and ventral lateral geniculate neurons to encode interocular differences in irradiance," *The Journal of physiology*, vol. 596, no. 22, pp. 5461–5481, 2018.

[33] K. Funke and F. Wrgötter, "On the significance of temporally structured activity in the dorsal lateral geniculate nucleus (lgn)," *Progress in neurobiology*, vol. 53, no. 1, pp. 67–119, 1997.

[34] A. Castonguay, S. Thomas, F. Lesage, and C. Casanova, "Repetitive and retinotopically restricted activation of the dorsal lateral geniculate nucleus with optogenetics," *PloS one*, vol. 9, no. 4, p. e94633, 2014.

[35] R. E. Kalil and M. Behan, "Synaptic reorganization in the dorsal lateral geniculate nucleus following damage to visual cortex in newborn or adult cats," *Journal of Comparative Neurology*, vol. 257, no. 2, pp. 216–236, 1987.

[36] F. Scala, D. Kobak, S. Shan, Y. Bernaerts, S. Laturnus, C. R. Cadwell, L. Hartmanis, E. Froudarakis, J. R. Castro, Z. H. Tan, *et al.*, "Layer 4 of mouse neocortex differs in cell types and circuit organization between sensory areas," *Nature communications*, vol. 10, no. 1, pp. 1–12, 2019.

[37] A. B. Butler and W. Hodos, *Comparative vertebrate neuroanatomy: evolution and adaptation*. John Wiley & Sons, 2005.

[38] C. E. Hagan, B. Bolon, and C. D. Keene, "Nervous system," in *Comparative Anatomy and Histology*, pp. 339–394, Elsevier, 2012.

[39] M. I. Sereno and J. Allman, "Cortical visual areas in mammals," *The neural basis of visual function*, vol. 4, pp. 160–172, 1991.

[40] C. B. Canto and M. P. Witter, "Cellular properties of principal neurons in the rat entorhinal cortex. i. the lateral entorhinal cortex," *Hippocampus*, vol. 22, no. 6, pp. 1256–1276, 2012.

[41] D. E. Haines and J. Lynch, *Fundamental Neuroscience for Basic and Clinical Applications*, ch. The Cerebral Cortex. Elsevier Health Sciences, 2012.

[42] D. Felleman, "Visual system in the brain," 2001.

[43] N. J. Priebe and D. Ferster, "Vision: Mechanisms of orientation, direction and depth," in *Encyclopedia of Neuroscience*, pp. 217–225, Elsevier Ltd, 2010.

[44] M. Ghodrati, S.-M. Khaligh-Razavi, and S. R. Lehky, "Towards building a more complex view of the lateral geniculate nucleus: recent advances in understanding its role," *Progress in neurobiology*, vol. 156, pp. 214–255, 2017.

[45] W. Ji, R. Gămănuţ, P. Bista, R. D. D'Souza, Q. Wang, and A. Burkhalter, "Modularity in the organization of mouse primary visual cortex," *Neuron*, vol. 87, no. 3, pp. 632–643, 2015.

[46] C. D. Gilbert and W. Li, "Top-down influences on visual processing," *Nature Reviews Neuroscience*, vol. 14, no. 5, pp. 350–363, 2013.

[47] S. Kondo, T. Yoshida, and K. Ohki, "Mixed functional microarchitectures for orientation selectivity in the mouse primary visual cortex," *Nature communications*, vol. 7, p. 13210, 2016.

[48] S. Schuett, T. Bonhoeffer, and M. Hübener, "Mapping retinotopic structure in mouse visual cortex with optical imaging," *Journal of Neuroscience*, vol. 22, no. 15, pp. 6549–6559, 2002.

[49] Blausen.com staff, 'Medical gallery of Blausen Medical 2014'. *WikiJournal of Medicine*, vol. 1, 2014, doi: `https://doi.org/10.15347/wjm/2014.010`, ISSN 2002-4436, [Accessed 10-May-2020].

[50] P. Dayan and L. Abbott, "Neural encoding ii: Reverse correlation and visual receptive fields," *Theoretical Neuroscience*, pp. 62–80, 2002.

[51] H. K. Hartline, "The response of single optic nerve fibers of the vertebrate eye to illumination of the retina," *American Journal of Physiology-Legacy Content*, vol. 121, no. 2, pp. 400–415, 1938.

[52] D. M. Piscopo, R. N. El-Danaf, A. D. Huberman, and C. M. Niell, "Diverse visual features encoded in mouse lateral geniculate nucleus," *Journal of Neuroscience*, vol. 33, no. 11, pp. 4642–4656, 2013.

[53] D. H. Hubel and T. N. Wiesel, "Receptive fields, binocular interaction and functional architecture in the cat's visual cortex," *The Journal of physiology*, vol. 160, no. 1, pp. 106–154, 1962.

[54] A. D. Lien and M. Scanziani, "Cortical direction selectivity emerges at convergence of thalamic synapses," *Nature*, vol. 558, no. 7708, pp. 80–86, 2018.

[55] All cell models used in the Allen V1 GLIF network are available from: `https://www.dropbox.com/sh/xb7xasih3d8027u/AADZTSNCvVM50bHZiuC_29sca/GLIF_network/models/cell_models?dl=0&subfolder_nav_tracking=1` [Accessed 12-May-2020].

[56] N. W. Gouwens, J. Berg, D. Feng, S. A. Sorensen, H. Zeng, M. J. Hawrylycz, C. Koch, and A. Arkhipov, "Systematic generation of biophysically detailed models for diverse cortical neuron types," *Nature communications*, vol. 9, no. 1, pp. 1–13, 2018.

[57] L. Paninski, J. Pillow, and J. Lewi, "Statistical models for neural encoding, decoding, and optimal stimulus design," *Progress in brain research*, vol. 165, pp. 493–507, 2007.

[58] N. W. Schultheiss, A. A. Prinz, and R. J. Butera, *Phase response curves in neuroscience: theory, experiment, and analysis*. Springer Science & Business Media, 2011.

[59] M. A. Triplett and G. J. Goodhill, "Probabilistic encoding models for multivariate neural data," *Frontiers in neural circuits*, vol. 13, p. 1, 2019.

[60] E. Chichilnisky, "A simple white noise analysis of neuronal light responses," *Network: Computation in Neural Systems*, vol. 12, no. 2, pp. 199–213, 2001.

[61] V. Nair and G. E. Hinton, "Rectified linear units improve restricted boltzmann machines," in *Proceedings of the 27th international conference on machine learning (ICML-10)*, pp. 807–814, 2010.

[62] H. Ozeki, O. Sadakane, T. Akasaki, T. Naito, S. Shimegi, and H. Sato, "Relationship between excitation and inhibition underlying size tuning and contextual response modulation in the cat primary visual cortex," *Journal of Neuroscience*, vol. 24, no. 6, pp. 1428–1438, 2004.

[63] H. J. Alitto and W. M. Usrey, "Surround suppression and temporal processing of visual signals," *Journal of neurophysiology*, vol. 113, no. 7, pp. 2605–2617, 2015.

[64] M. W. Self, J. A. Lorteije, J. Vangeneugden, E. H. van Beest, M. E. Grigore, C. N. Levelt, J. A. Heimel, and P. R. Roelfsema, "Orientation-tuned surround suppression in mouse visual cortex," *Journal of Neuroscience*, vol. 34, no. 28, pp. 9290–9304, 2014.

[65] T. G. Fisher, H. J. Alitto, and W. M. Usrey, "Retinal and nonretinal contributions to extraclassical surround suppression in the lateral geniculate nucleus," *Journal of Neuroscience*, vol. 37, no. 1, pp. 226–235, 2017.

[66] H. J. Alitto and W. M. Usrey, "Origin and dynamics of extraclassical suppression in the lateral geniculate nucleus of the macaque monkey," *Neuron*, vol. 57, no. 1, pp. 135–146, 2008.

[67] M. Zhu and C. J. Rozell, "Visual nonclassical receptive field effects emerge from sparse coding in a dynamical system," *PLoS computational biology*, vol. 9, no. 8, 2013.

[68] I. M. Gelfand, *Lectures on linear algebra*. Courier Corporation, 1989.

[69] M.-P. Schallmo and S. O. Murray, "Identifying separate components of surround suppression," *Journal of vision*, vol. 16, no. 1, pp. 2–2, 2016.

[70] V. Bonin, V. Mante, and M. Carandini, "The suppressive field of neurons in lateral geniculate nucleus," *Journal of Neuroscience*, vol. 25, no. 47, pp. 10844–10856, 2005.

[71] S. G. Solomon, A. J. White, and P. R. Martin, "Extraclassical receptive field properties of parvocellular, magnocellular, and koniocellular cells in the primate lateral geniculate nucleus," *Journal of Neuroscience*, vol. 22, no. 1, pp. 338–349, 2002.

[72] J. Allman, F. Miezin, and E. McGuinness, "Stimulus specific responses from beyond the classical receptive field: neurophysiological mechanisms for local-global comparisons in visual neurons," *Annual review of neuroscience*, vol. 8, no. 1, pp. 407–430, 1985.

[73] J. R. Cavanaugh, W. Bair, and J. A. Movshon, "Nature and interaction of signals from the receptive field center and surround in macaque v1 neurons," *Journal of neurophysiology*, vol. 88, no. 5, pp. 2530–2546, 2002.

[74] C. A. Henry, S. Joshi, D. Xing, R. M. Shapley, and M. J. Hawken, "Functional characterization of the extraclassical receptive field in macaque v1: contrast, orientation, and temporal dynamics," *Journal of Neuroscience*, vol. 33, no. 14, pp. 6230–6242, 2013.

[75] M. A. Smith, "Surround suppression in the early visual system," *Journal of Neuroscience*, vol. 26, no. 14, pp. 3624–3625, 2006.

[76] I. M. Andolina, H. E. Jones, and A. M. Sillito, "Effects of cortical feedback on the spatial properties of relay cells in the lateral geniculate nucleus," *Journal of neurophysiology*, vol. 109, no. 3, pp. 889–899, 2013.

[77] H. E. Jones, I. M. Andolina, B. Ahmed, S. D. Shipp, J. T. Clements, K. L. Grieve, J. Cudeiro, T. E. Salt, and A. M. Sillito, "Differential feedback modulation of center and surround mechanisms in parvocellular cells in the visual thalamus," *Journal of Neuroscience*, vol. 32, no. 45, pp. 15946–15951, 2012.

[78] M. J. Nolt, R. D. Kumbhani, and L. A. Palmer, "Suppression at high spatial frequencies in the lateral geniculate nucleus of the cat," *Journal of neurophysiology*, vol. 98, no. 3, pp. 1167–1180, 2007.

[79] P. Murphy and A. Sillito, "Corticofugal feedback influences the generation of length tuning in the visual pathway," *Nature*, vol. 329, no. 6141, pp. 727–729, 1987.

[80] S. R. Olsen, D. S. Bortone, H. Adesnik, and M. Scanziani, "Gain control by layer six in cortical circuits of vision," *Nature*, vol. 483, no. 7387, pp. 47–52, 2012.

[81] B. S. Webb, C. J. Tinsley, C. J. Vincent, and A. M. Derrington, "Spatial distribution of suppressive signals outside the classical receptive field in lateral geniculate nucleus," *Journal of neurophysiology*, vol. 94, no. 3, pp. 1789–1797, 2005.

[82] A. Vaiceliunaite, S. Erisken, F. Franzen, S. Katzner, and L. Busse, "Spatial integration in mouse primary visual cortex," *Journal of neurophysiology*, vol. 110, no. 4, pp. 964–972, 2013.

[83] W. W. Tschetter, G. Govindaiah, I. M. Etherington, W. Guido, and C. M. Niell, "Refinement of spatial receptive fields in the developing mouse lateral geniculate nucleus is coordinated with excitatory and inhibitory remodeling," *Journal of Neuroscience*, vol. 38, no. 19, pp. 4531–4542, 2018.

[84] M. Fahle, T. Poggio, T. A. Poggio, *et al.*, *Perceptual learning*. MIT Press, 2002.

[85] R. Shapley and P. Lennie, "Spatial frequency analysis in the visual system," *Annual review of neuroscience*, vol. 8, no. 1, pp. 547–581, 1985.

[86] B. Cleland, T. Harding, and U. Tulunay-Keesey, "Visual resolution and receptive field size: examination of two kinds of cat retinal ganglion cell," *Science*, vol. 205, no. 4410, pp. 1015–1017, 1979.

[87] C. Enroth-Cugell and J. G. Robson, "Functional characteristics and diversity of cat retinal ganglion cells. basic characteristics and quantitative description.," *Investigative ophthalmology & visual science*, vol. 25, no. 3, pp. 250–267, 1984.

[88] A. Adolph, "Center-surround, orientation, and directional properties of turtle retinal horizontal cells," *Biological cybernetics*, vol. 58, no. 6, pp. 373–385, 1988.

[89] D. A. Pollen and J. P. Gaska, "Vision, visual cortex, and frequency analysis," in *Sensory System I*, pp. 69–70, Springer, 1988.

[90] P. M. DiLorenzo and J. D. Victor, *Spike timing: mechanisms and function*. CRC Press, 2013.

[91] H. Plesser and G. Einevoll, "Linear mechanistic models for the dorsal lateral geniculate nucleus of cat probed using drifting-grating stimuli," *Network: Computation in neural systems*, vol. 13, no. 4, pp. 503–530, 2002.

[92] J. W. Pillow, L. Paninski, V. J. Uzzell, E. P. Simoncelli, and E. Chichilnisky, "Prediction and decoding of retinal ganglion cell responses with a probabilistic spiking model," *Journal of Neuroscience*, vol. 25, no. 47, pp. 11003–11013, 2005.

[93] R. W. Rodieck, "Quantitative analysis of cat retinal ganglion cell response to visual stimuli," *Vision research*, vol. 5, no. 12, pp. 583–601, 1965.

[94] D. Gabor, "Theory of communication. part 1: The analysis of information," *Journal of the Institution of Electrical Engineers-Part III: Radio and Communication Engineering*, vol. 93, no. 26, pp. 429–441, 1946.

[95] G. T. Einevoll and H. E. Plesser, "Response of the difference-of-gaussians model to circular drifting-grating patches," *Visual neuroscience*, vol. 22, no. 4, pp. 437–446, 2005.

[96] G. Buračas, ". mathematical modelling in the early visual system: Why and how?," *Modulation of neuronal responses: implications for active vision*, vol. 334, p. 135, 2003.

[97] J. G. Daugman, "Two-dimensional spectral analysis of cortical receptive field profiles," *Vision research*, vol. 20, no. 10, pp. 847–856, 1980.

[98] J. G. Daugman, "Uncertainty relation for resolution in space, spatial frequency, and orientation optimized by two-dimensional visual cortical filters," *JOSA A*, vol. 2, no. 7, pp. 1160–1169, 1985.

[99] S. Marĉelja, "Mathematical description of the responses of simple cortical cells," *JOSA*, vol. 70, no. 11, pp. 1297–1300, 1980.

[100] D. J. Field and D. J. Tolhurst, "The structure and symmetry of simple-cell receptive-field profiles in the cat's visual cortex," *Proceedings of the Royal society of London. Series B. Biological sciences*, vol. 228, no. 1253, pp. 379–400, 1986.

[101] T. Lindeberg, "A computational theory of visual receptive fields," *Biological cybernetics*, vol. 107, no. 6, pp. 589–635, 2013.

[102] J. P. Jones and L. A. Palmer, "An evaluation of the two-dimensional gabor filter model of simple receptive fields in cat striate cortex," *Journal of neurophysiology*, vol. 58, no. 6, pp. 1233–1258, 1987.

[103] D. Cope, B. Blakeslee, and M. E. McCourt, "Simple cell response properties imply receptive field structure: balanced gabor and/or bandlimited field functions," *JOSA A*, vol. 26, no. 9, pp. 2067–2092, 2009.

[104] M. P. Sceniak, M. J. Hawken, and R. Shapley, "Contrast-dependent changes in spatial frequency tuning of macaque v1 neurons: effects of a changing receptive field size," *Journal of Neurophysiology*, vol. 88, no. 3, pp. 1363–1373, 2002.

[105] R. A. Young, "The gaussian derivative model for spatial vision: I. retinal mechanisms," *Spatial vision*, vol. 2, no. 4, pp. 273–293, 1987.

[106] Code and data files for the Allen Institute Mouse Primary Visual Cortex Model: `https://www.dropbox.com/sh/w5u31m3hq6u2x5m/AACpYpeWnm6s_qJDpmgrYgP7a?dl=0`, [Accessed 10-March-2020].

[107] C. Koch, C. Reid, H. Zeng, S. Mihalas, M. Hawrylycz, J. Philips, C. Dang, and A. Jones, "Project mindscope," in *The Future of the Brain: Essays by the World's Leading Neuroscientists*, p. 31, Princeton University Press, 2014.

[108] B. Tasic, V. Menon, T. N. Nguyen, T. K. Kim, T. Jarsky, Z. Yao, B. Levi, L. T. Gray, S. A. Sorensen, T. Dolbeare, *et al.*, "Adult mouse cortical cell taxonomy revealed by single cell transcriptomics," *Nature neuroscience*, vol. 19, no. 2, p. 335, 2016.

[109] B. Rudy, G. Fishell, S. Lee, and J. Hjerling-Leffler, "Three groups of interneurons account for nearly 100% of neocortical gabaergic neurons," *Developmental neurobiology*, vol. 71, no. 1, pp. 45–61, 2011.

[110] X.-y. Ji, B. Zingg, L. Mesik, Z. Xiao, L. I. Zhang, and H. W. Tao, "Thalamocortical innervation pattern in mouse auditory and visual cortex: laminar and cell-type specificity," *Cerebral Cortex*, vol. 26, no. 6, pp. 2612–2625, 2016.

[111] Script defining the main functions of the LGN model, `lgn_functions.py`, available from: `https://www.dropbox.com/sh/w5u31m3hq6u2x5m/AABcUXYv2HPbUYldtXreL6Gxa/LGN/scripts?dl=0&preview=lgn_functions.py&subfolder_nav_tracking=1` [Accessed 29-April-2020].

[112] The file `lgn_full_col_cells_3.csv`, which lists the parameters for each LGN cell, is available from: `https://www.dropbox.com/sh/xb7xasih3d8027u/AADqmH2hSsp8sd_F7EWmT8Kka/LGN/LGN?dl=0&preview=lgn_full_col_cells_3.csv&subfolder_nav_tracking=1` [Accessed 01-May-2020].

[113] M. S. Grubb and I. D. Thompson, "Quantitative characterization of visual response properties in the mouse dorsal lateral geniculate nucleus," *Journal of neurophysiology*, vol. 90, no. 6, pp. 3594–3607, 2003.

[114] S. J. Jackson, N. Andrews, D. Ball, I. Bellantuono, J. Gray, L. Hachoumi, A. Holmes, J. Latcham, A. Petrie, P. Potter, *et al.*, "Does age matter? the impact of rodent age on study outcomes," *Laboratory animals*, vol. 51, no. 2, pp. 160–169, 2017.

[115] H. Adesnik, W. Bruns, H. Taniguchi, Z. J. Huang, and M. Scanziani, "A neural circuit for spatial summation in visual cortex," *Nature*, vol. 490, no. 7419, pp. 226–231, 2012.

[116] G. Van den Bergh, B. Zhang, L. Arckens, and Y. M. Chino, "Receptive-field properties of v1 and v2 neurons in mice and macaque monkeys," *Journal of Comparative Neurology*, vol. 518, no. 11, pp. 2051–2070, 2010.

[117] U. C. Dräger, "Receptive fields of single cells and topography in mouse visual cortex," *Journal of Comparative Neurology*, vol. 160, no. 3, pp. 269–289, 1975.

[118] J. H. Marshel, M. E. Garrett, I. Nauhaus, and E. M. Callaway, "Functional specialization of seven mouse visual cortical areas," *Neuron*, vol. 72, no. 6, pp. 1040–1054, 2011.

[119] J. Zhuang, Y. Bereshpolova, C. R. Stoelzel, J. M. Huff, X. Hei, J.-M. Alonso, and H. A. Swadlow, "Brain state effects on layer 4 of the awake visual cortex," *Journal of Neuroscience*, vol. 34, no. 11, pp. 3888–3900, 2014.

[120] M. L. Andermann, A. M. Kerlin, D. K. Roumis, L. L. Glickfeld, and R. C. Reid, "Functional specialization of mouse higher visual cortical areas," *Neuron*, vol. 72, no. 6, pp. 1025–1039, 2011.

[121] E. LeDue, M. Zou, and N. Crowder, "Spatiotemporal tuning in mouse primary visual cortex," *Neuroscience letters*, vol. 528, no. 2, pp. 165–169, 2012.

[122] E. Gao, G. C. DeAngelis, and A. Burkhalter, "Parallel input channels to mouse primary visual cortex," *Journal of Neuroscience*, vol. 30, no. 17, pp. 5912–5926, 2010.

[123] C. M. Niell and M. P. Stryker, "Highly selective receptive fields in mouse visual cortex," *Journal of Neuroscience*, vol. 28, no. 30, pp. 7520–7536, 2008.

[124] T. Bezdudnaya, M. Cano, Y. Bereshpolova, C. R. Stoelzel, J.-M. Alonso, and H. A. Swadlow, "Thalamic burst mode and inattention in the awake lgnd," *Neuron*, vol. 49, no. 3, pp. 421–432, 2006.

[125] D. B. Rubin, S. D. Van Hooser, and K. D. Miller, "The stabilized supralinear network: a unifying circuit motif underlying multi-input integration in sensory cortex," *Neuron*, vol. 85, no. 2, pp. 402–417, 2015.

[126] File specifying the recurrent connections in the V1 model, `v1_v1_edges.h5`, available from: `https://www.dropbox.com/sh/w5u31m3hq6u2x5m/AACDWH47Jmmo-kHcTJibpnPOa/GLIF_network/network?dl=0&preview=v1_v1_edges.h5&subfolder_nav_tracking=1` [Accessed 01-May-2020].

[127] File specifying parameters for recurrent connections in the V1 model, `v1_v1_edge_types.csv`, available from: `https://www.dropbox.com/sh/w5u31m3hq6u2x5m/AACDWH47Jmmo-kHcTJibpnPOa/GLIF_network/network?dl=0&preview=v1_v1_edge_types.csv&subfolder_nav_tracking=1` [Accessed 01-May-2020].

[128] Configuration file for running the V1 model, `config.json`, available from: `https://www.dropbox.com/sh/w5u31m3hq6u2x5m/AADiliDyiS8BIxaDgWXS6pwQa/GLIF_network?dl=0&preview=config.json&subfolder_nav_tracking=1` [Accessed 01-May-2020].

[129] M. Carandini and D. J. Heeger, "Normalization as a canonical neural computation," *Nature Reviews Neuroscience*, vol. 13, no. 1, p. 51, 2012.

[130] The Brain Modeling Toolkit (BMTK), available from: `www.github.com/AllenInstitute/bmtk/blob/develop/bmtk/simulator/utils/nwb.py` [Accessed 29-April-2020].

[131] Script for running the LGN model, available from: `https://www.dropbox.com/sh/xb7xasih3d8027u/AAArO8IBUk-ovTsVUuoG0MjNa/LGN/scripts?dl=0&preview=simulate_drifiting_gratings.py&subfolder_nav_tracking=1` [Accessed 01-May-2020].

[132] Script for running the V1 model, available from: `https://www.dropbox.com/sh/xb7xasih3d8027u/AAB38AiQLeRViJpyZEA83wgTa/GLIF_network?dl=0&preview=run_pointnet.py&subfolder_nav_tracking=1` [Accessed 01-May-2020].

[133] Script for creating full-field drifting grating stimuli, available from: `https://www.dropbox.com/sh/xb7xasih3d8027u/AAADViBKH980roqFqdtE1Lw6a/LGN/scripts/lgnmodel/lgnmodel?dl=0&preview=movie.py&subfolder_nav_tracking=1` [Accessed 01-May-2020].

[134] J. A. Movshon, I. D. Thompson, and D. J. Tolhurst, "Spatial summation in the receptive fields of simple cells in the cat's striate cortex.," *The Journal of physiology*, vol. 283, no. 1, pp. 53–77, 1978.

[135] J. A. Movshon, I. D. Thompson, and D. J. Tolhurst, "Receptive field organization of complex cells in the cat's striate cortex.," *The Journal of physiology*, vol. 283, no. 1, pp. 79–99, 1978.

[136] J. W. Cooley and J. W. Tukey, "An algorithm for the machine calculation of complex fourier series," *Mathematics of computation*, vol. 19, no. 90, pp. 297–301, 1965.

[137] Documentation for the Discrete Fourier Transform implemented in Python, via numpy.fft: `https://numpy.org/doc/1.18/reference/routines.fft.html#module-numpy.fft` [Accessed 26-April-2020].

[138] S. A. Teukolsky, B. P. Flannery, W. Press, and W. Vetterling, "Numerical recipes in c," *SMR*, vol. 693, no. 1, pp. 59–70, 1992.

[139] The file `GLIF_network/network/v1_nodes.h5`, which lists the neuron coordinates for the V1 model, is available from: `https://www.dropbox.com/sh/w5u31m3hq6u2x5m/AACDWH47Jmmo-kHcTJibpnPOa/GLIF_network/network?dl=0&preview=v1_nodes.h5&subfolder_nav_tracking=1` [Accessed 01-May-2020].

[140] C. Koch, M. S. Gazzaniga, T. F. Heatherton, J. E. Ledoux, and N. Logothetis, *The Cognitive Neurosciences*. MIT Press, 2004.

[141] C. D. Gilbert and T. N. Wiesel, "Clustered intrinsic connections in cat visual cortex," *Journal of Neuroscience*, vol. 3, no. 5, pp. 1116–1133, 1983.

[142] D. J. Heeger, "Normalization of cell responses in cat striate cortex," *Visual neuroscience*, vol. 9, no. 2, pp. 181–197, 1992.

[143] J. Nelson and B. Frost, "Orientation-selective inhibition from beyond the classic visual receptive field," *Brain research*, vol. 139, no. 2, pp. 359–365, 1978.

[144] C. D. Gilbert and T. N. Wiesel, "The influence of contextual stimuli on the orientation selectivity of cells in primary visual cortex of the cat," *Vision research*, vol. 30, no. 11, pp. 1689–1701, 1990.

[145] W. Bair, J. R. Cavanaugh, and J. A. Movshon, "Time course and time-distance relationships for surround suppression in macaque v1 neurons," *Journal of Neuroscience*, vol. 23, no. 20, pp. 7690–7701, 2003.

[146] M. H. Mobarhan, G. Halnes, T. Hafting, M. Fyhn, G. T. Einevoll, *et al.*, "Firing-rate based network modeling of the dlgn circuit: Effects of cortical feedback on spatiotemporal response properties of relay cells," *PLoS computational biology*, vol. 14, no. 5, p. e1006156, 2018.

[147] Y. Umino, E. Solessio, and R. B. Barlow, "Speed, spatial, and temporal tuning of rod and cone vision in mouse," *Journal of Neuroscience*, vol. 28, no. 1, pp. 189–198, 2008.

[148] List of all files specifying the connections to V1 cells, available from: `https://www.dropbox.com/sh/xb7xasih3d8027u/AACk20E6arDcY8OmSWAi2TGba/GLIF_network/network?dl=0&subfolder_nav_tracking=1` [Accessed 01-May-2020].

[149] V. Vaingankar, C. Soto Sanchez, X. Wang, F. T. Sommer, and J. A. Hirsch, "Neurons in the thalamic reticular nucleus are selective for diverse and complex visual features," *Frontiers in integrative neuroscience*, vol. 6, p. 118, 2012.

[150] K. Rottmann, "Matematisk formelsamling, norsk utgave," *Spektrum Forlag*, vol. 15, 2015.

Appendix A

Derivations of Filter Responses

Response to full-field grating

The spatial frequency tuning of the DoG filter in equation (5.6) can be found from integral 5.2. The full-field drifting grating stimuli is described by equation (5.7), where the contrast, A, is set to unity. Substituting for the stimulus and spatial filter, the integral becomes

$$\begin{aligned} L(t) = &\int_{\tau=0}^{\tau=\infty} \int_{x=-\infty}^{x=\infty} \int_{y=-\infty}^{y=\infty} D_t(\tau) \cos(kx - \omega(t-\tau)) \\ &\times \left[\frac{A_c}{2\pi\sigma_c^2} e^{-(x^2+y^2)/2\sigma_c^2} - \frac{A_s}{2\pi\sigma_s^2} e^{-(x^2+y^2)/2\sigma_s^2} \right] d\tau dx dy. \end{aligned} \tag{A.1}$$

Expressing $\cos(kx - \omega(t-\tau))$ as $Re\left\{e^{i(kx-\omega(t-\tau))}\right\}$ and extracting the terms independent of the spatial coordinates, this can be simplified to

$$\begin{aligned} L(t) = &Re\Big\{ \Lambda(\omega) e^{-i\omega t} \int_{x=-\infty}^{x=\infty} \int_{y=-\infty}^{y=\infty} \left[\frac{A_c}{2\pi\sigma_c^2} e^{-(x^2+y^2)/2\sigma_c^2} \right. \\ &\left. - \frac{A_s}{2\pi\sigma_s^2} e^{-(x^2+y^2)/2\sigma_s^2} \right] e^{ikx} dx dy \Big\} \end{aligned} \tag{A.2}$$

where $\Lambda(\omega)$ represents the integral over τ, which is the contribution of the temporal kernel. The integral over y is a standard Gaussian integral [150, p. 155], giving

$$\begin{aligned} L(t) = Re\Big\{ \Lambda(\omega) e^{-iwt} \int_{x=-\infty}^{x=\infty} \left[\frac{A_c}{\sqrt{2\pi\sigma_c^2}} e^{-x^2/2\sigma_c^2} \right. \\ \left. - \frac{A_s}{\sqrt{2\pi\sigma_s^2}} e^{-x^2/2\sigma_s^2} \right] e^{ikx} dx \Big\}. \end{aligned} \tag{A.3}$$

Furthermore, $\Lambda(\omega)$ can be written as a complex number, $|\Lambda(\omega)|\, e^{i\theta(\omega)}$, with magnitude $|\Lambda(\omega)|$ and phase $\theta(\omega)$. The remaining integral over x can be solved using a Fourier

transformation [150, p. 168], which gives the result

$$\begin{aligned} L(t) = Re\Big\{ \; |\Lambda(\omega)|\, e^{i\theta(\omega)} e^{-i\omega t} \Big\} \Big[A_c e^{-k^2\sigma_c^2/2} \\ - \; A_s e^{-k^2\sigma_s^2/2} \Big]. \\ = |\Lambda(\omega)| \cos(\omega t - \theta(\omega)) \Big[A_c e^{-k^2\sigma_c^2/2} \\ - \; A_s e^{-k^2\sigma_s^2/2} \Big] \\ = L_t(t) L_s. \end{aligned} \tag{A.4}$$

This equation describes the response of the DoG filter to the grating stimuli. The spatial frequency tuning is defined by the term in square brackets. The temporal dependency is an oscillatory cos-term depending on the temporal frequency of the grating. For $A_s = 0$, the DoG filter reduces to a single Gaussian filter, and the response becomes

$$L(t) = |\Lambda(\omega)|\, A_c e^{-k^2\sigma_c^2/2} \cos(\omega t - \theta(\omega)). \tag{A.5}$$

The spatial frequency tuning of equation (A.4) and (A.5) correspond to L_s in equation (5.9) and 5.8 respectively.

Response to an isoluminous spot

The response of a DoG filter to an isoluminous spot can be derived from the integral 5.2. The limits on the integral will now define the radius of the circle. For maximum luminance, the stimulus $S(x, y, t-\tau)$ is equal to one. Substituting the DoG filter 5.6 for the spatial kernel and changing to polar coordinates, one gets

$$\begin{aligned} L(t) =& 2\pi \int_0^\infty D_t(\tau) d\tau \int_{r=0}^{r=a} \left[\frac{A_c}{2\pi\sigma_c^2} e^{-r^2/2\sigma_c^2} - \frac{A_s}{2\pi\sigma_s^2} e^{-r^2/2\sigma_s^2} \right] r dr \\ \propto & \int_{r=0}^{r=a} \left[\frac{A_c}{\sigma_c^2} e^{-r^2/2\sigma_c^2} - \frac{A_s}{\sigma_s^2} e^{-r^2/2\sigma_s^2} \right] r dr \end{aligned} \tag{A.6}$$

where a is the radius of the spot. The second integral can be solved by using the substitution $u = r^2$, giving

$$L(t) \propto \left[A_c(1 - e^{-a^2/2\sigma_c^2}) - A_s(1 - e^{-a^2/2\sigma_s^2}) \right] \tag{A.7}$$

which is the equation (5.12). A single Gaussian spatial filter has $A_s = 0$, so that the result is reduced to

$$L(t) \propto \left[(1 - e^{-a^2/2\sigma_c^2}) \right], \tag{A.8}$$

that is, equation (5.11).

Appendix B

Figures

B.1 Original Allen Model

B.1.1 Spatial frequency tuning - full-field gratings

LGN

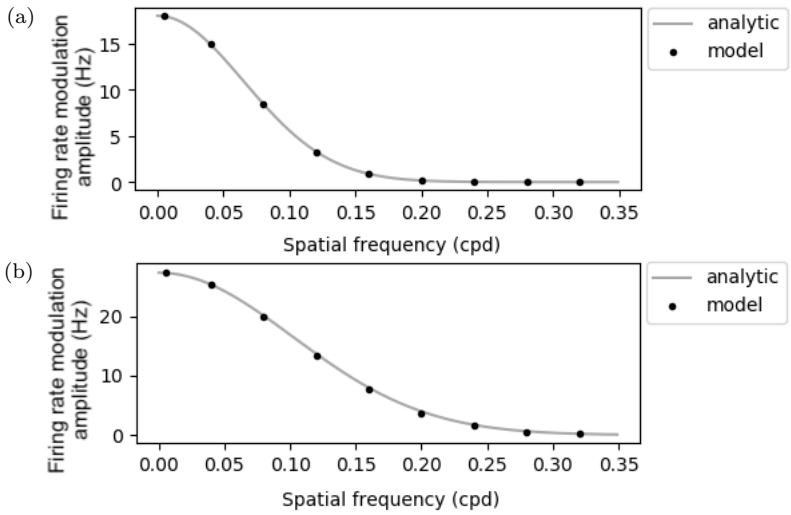

Figure B.1: Spatial frequency tuning of two neurons in the original Allen LGN model. The analytical solution from equation (5.8) (grey line) and simulated results (black dots) are compared. The neurons have receptive field sizes defined by equation (5.4), with $\sigma = 2.44^{\circ}$ (a) and $\sigma = 1.56^{\circ}$ (b). The neuron ids and classes are 24, sON-TF8 (a) and 2775, sON-TF4 (b).

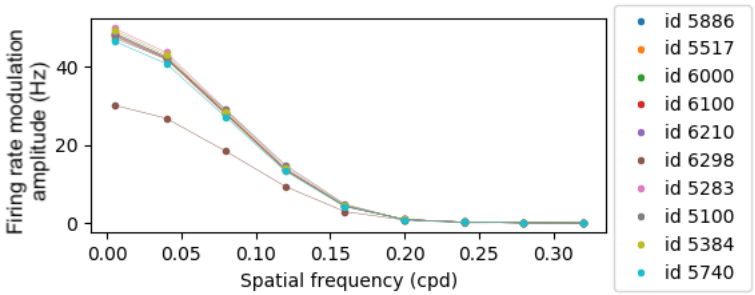

Figure B.2: Spatial frequency tuning curves of ten randomly selected sONsOFF neurons, obtained from LGN simulations. The sONsOFF have receptive fields consisting of two Gaussians filters (equation (5.4)) of opposite polarity, separated by a few degrees. Each subfilter has $\sigma = 2°$.

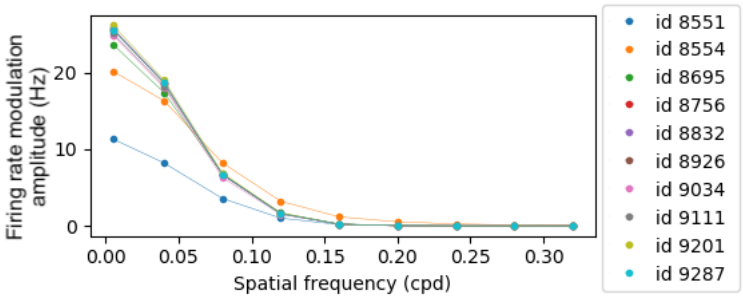

Figure B.3: Spatial frequency tuning curves of ten randomly selected sONtOFF neurons in the LGN, obtained from simulations. The sONtOFF have receptive fields consisting of two Gaussians filters (equation (5.4)) of opposite polarity, separated by a few degrees. Each subfilter has $\sigma = 3°$.

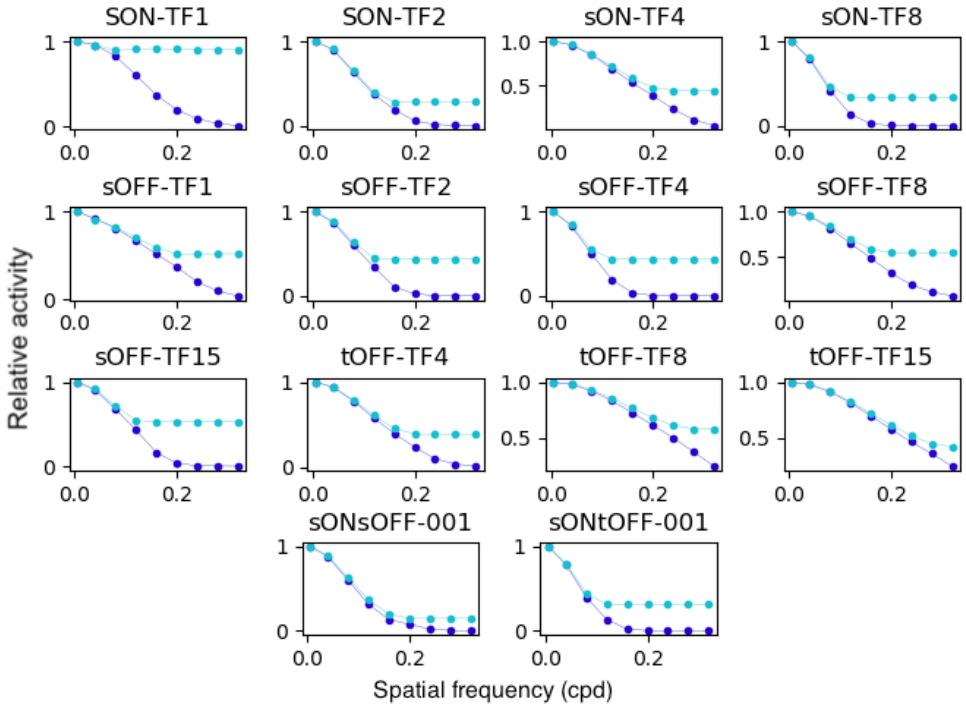

Figure B.4: Spatial frequency tuning of original LGN neurons, quantified by the normalised F0 (cyan) and F1 (dark blue) responses. The stimuli were full-field drifting gratings with a temporal frequency of 4 Hz.

V1

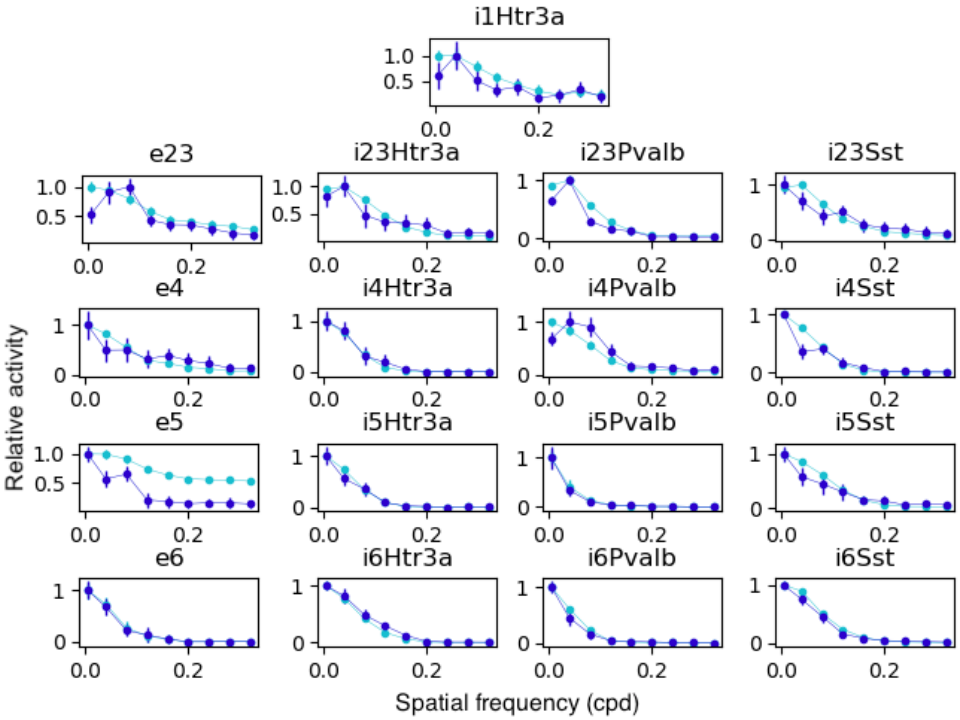

Figure B.5: Spatial frequency tuning of original V1 neurons, quantified by the normalised F0 (cyan) and F1 (dark blue) responses. The stimuli were full-field drifting gratings with a temporal frequency of 4 Hz. Curves show the mean over 20 trials of 3-second simulations, with error bars indicating the standard deviation.

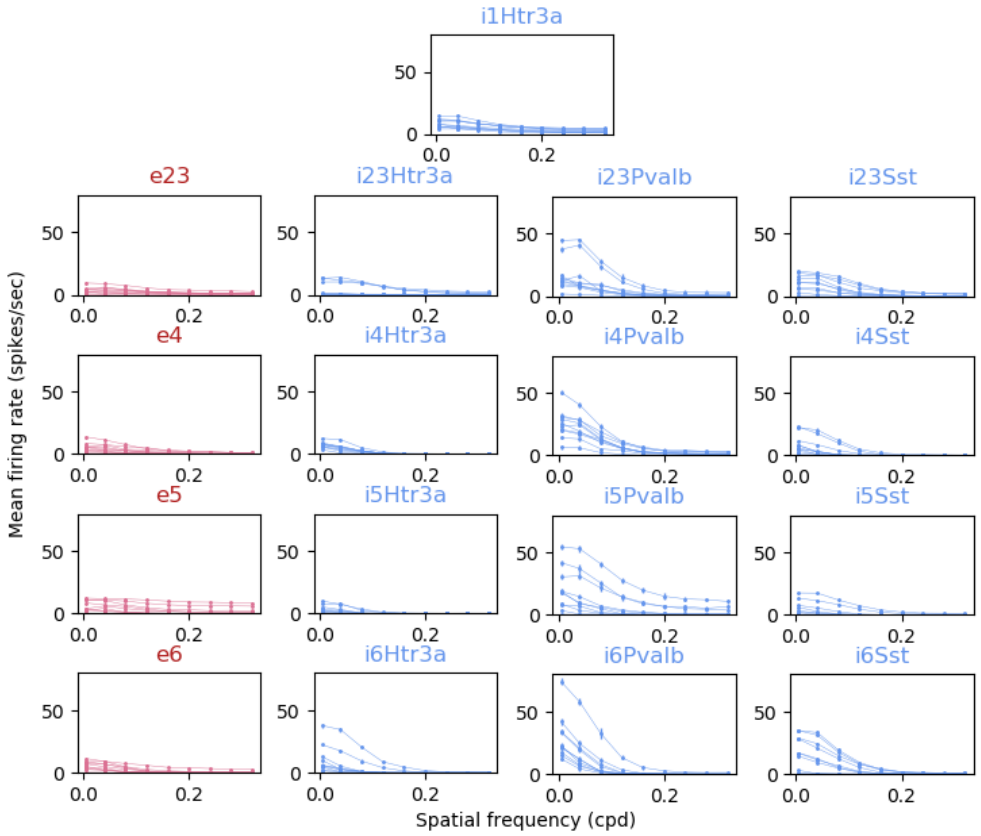

Figure B.6: Spatial frequency tuning of V1 neurons in the original Allen model. Each subplot shows 10 responsive neurons from the respective classes. The stimuli were a full-field drifting gratings with temporal frequency of 4 Hz. Curves show the mean over 20 trials of 3-second simulations, with error bars indicating the standard deviation.

B.1.2 Size tuning - flashing spots

V1

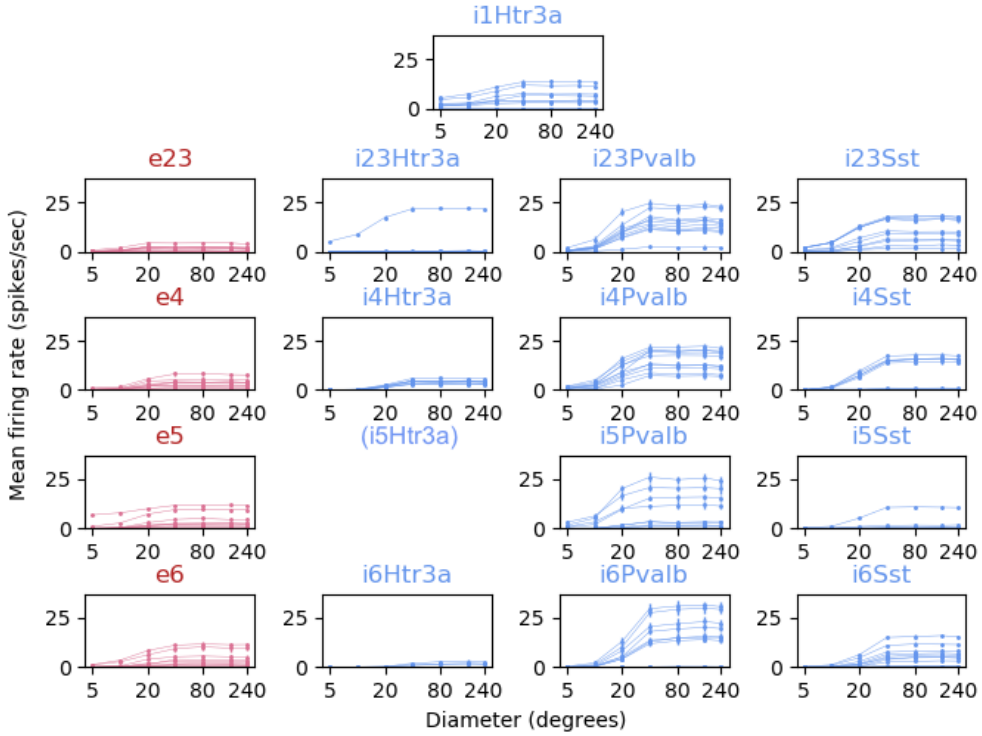

Figure B.7: Size tuning in the original V1 for white flashing spots of 75 ms duration. Up to ten neurons from each population are included, depending on the number of cells within 70 µm of stimuli center. For the i5Htr3a class, there were no cells within this radius. The measurements show mean responses over 20 trials, with error bars representing the standard deviation.

B.1.3 Size tuning - patch gratings

LGN

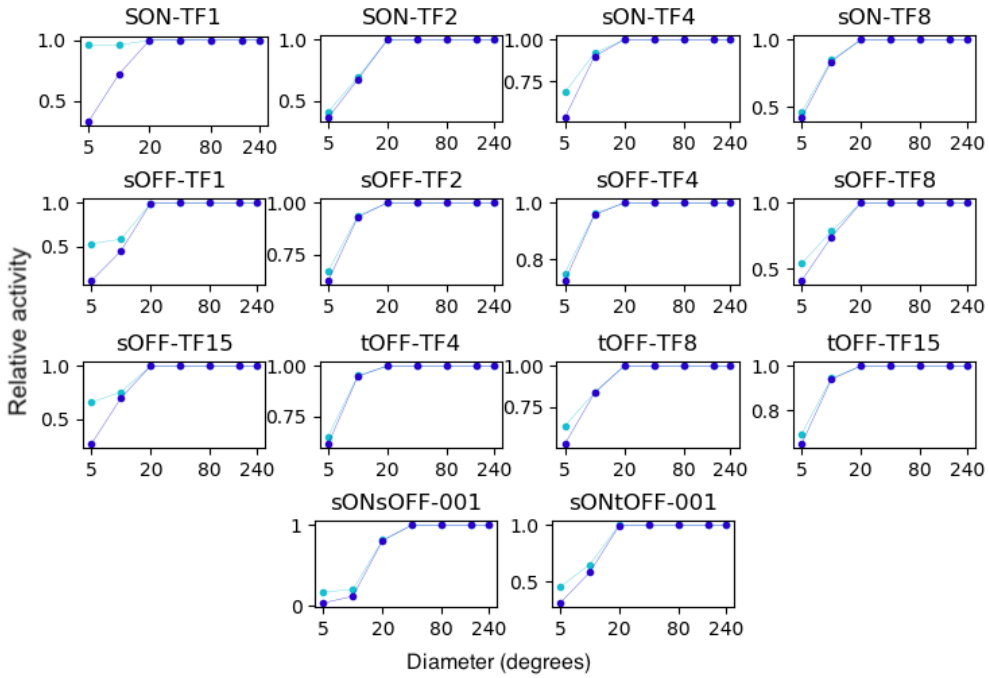

Figure B.8: Size tuning of original LGN neurons, quantified by the normalised F0 (cyan) and F1 (dark blue) responses. The stimuli were patches of drifting grating with temporal and spatial frequencies of 4 Hz and 0.04 cpd.

V1

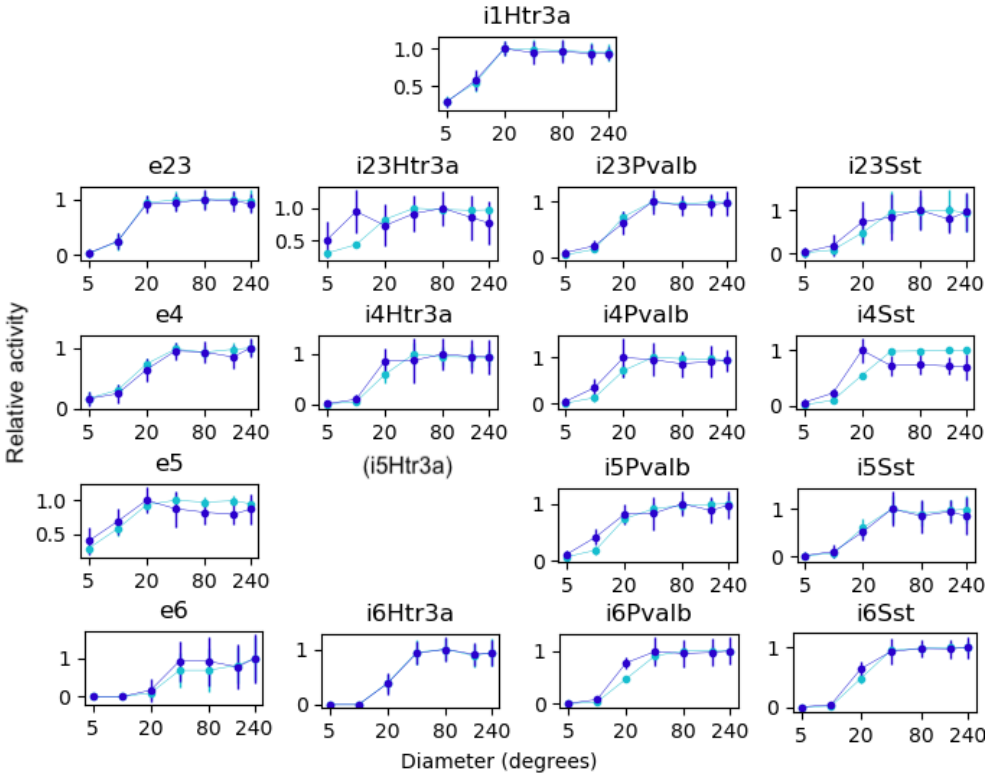

Figure B.9: Size tuning of original V1 neurons, quantified by the normalised F0 (cyan) and F1 (dark blue) responses. The stimuli were patch grating with spatial and temporal frequencies of 0.04 cpd and 4 Hz.

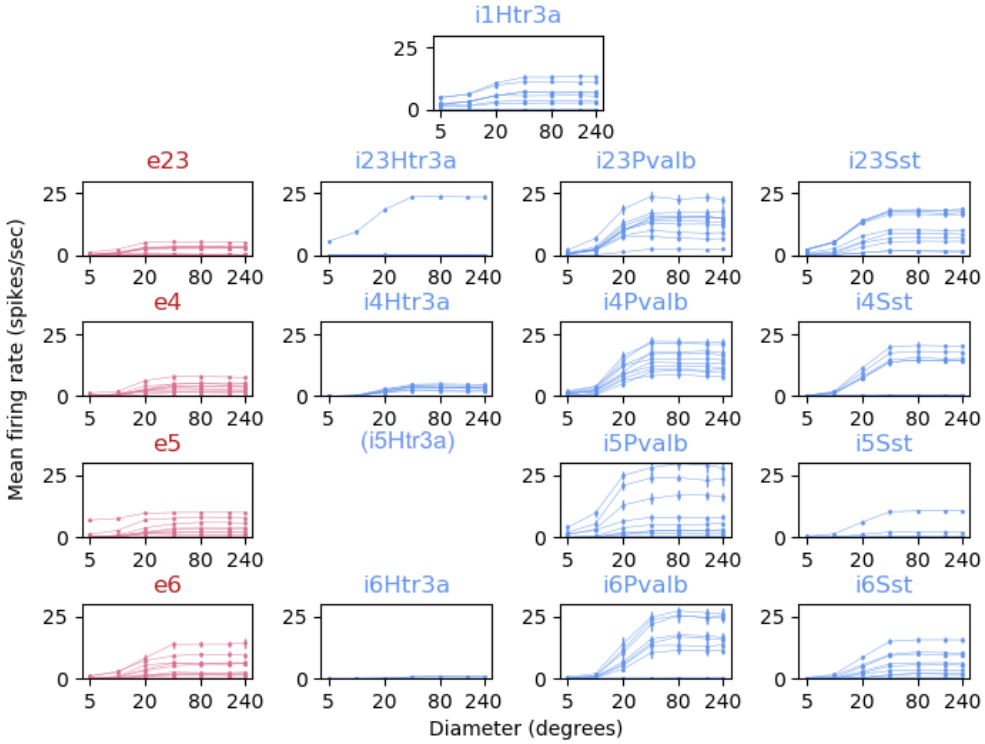

Figure B.10: Size tuning of V1 neurons for patch grating stimuli in the original Allen model. Up to ten neurons from each population are included, depending on the number of cells within 70µm of stimuli center. For the i5Htr3a class, there were no cells within this radius. Mean firing rate over 20 trials is plotted, with error bars showing the standard deviation. The patch gratings had spatial and temporal frequencies of 0.04 cpd and 4 Hz.

B.2 M1 Allen Model

B.2.1 Spatial frequency tuning - full-field gratings

LGN

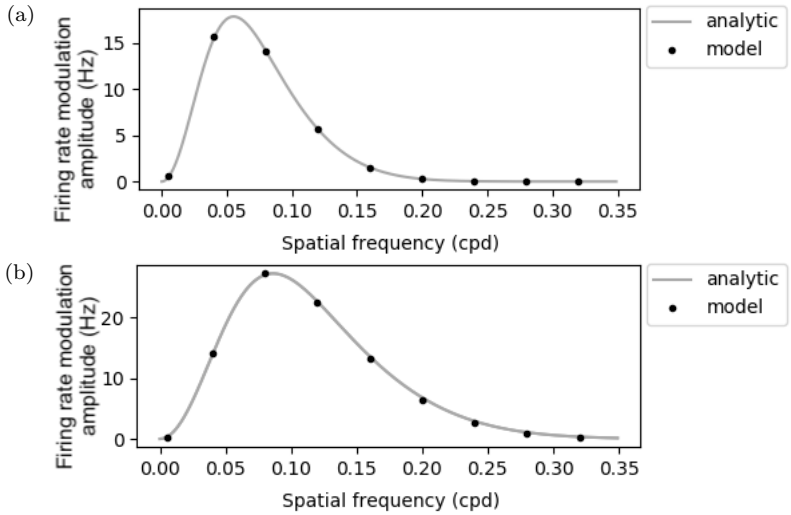

Figure B.11: Spatial frequency tuning for two neurons from the M1 LGN model. The analytical solution from equation (5.9) (grey line) and simulated results (black dots) are compared. The neurons have receptive fields defined by equation (5.6). (a) The neuron (id = 24, type = sON-TF8) has receptive field size with $\sigma_c = 2.44°$. (b) The neuron (id = 2775, type = sON-TF4) has a receptive field size with $\sigma_c = 1.56°$.

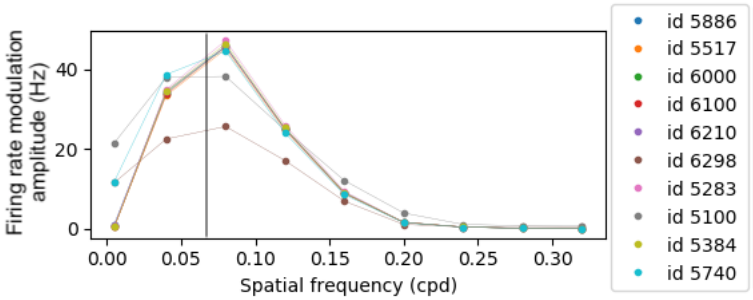

Figure B.12: Spatial frequency tuning curves of ten randomly selected sONsOFF neurons, obtained from LGN simulations. The vertical black line indicates the predicted preferred spatial frequency obtained when treating the neurons as a single ON or OFF cell with a circular receptive field. This value (0.067 cpd) is used as the estimated optimal frequency of the sONsOFF cells in histogram 9.10 section 9.2.3. The sONsOFF neurons have receptive fields consisting of two DoG filters (equation (5.6)) of opposite polarity, separated by a few degrees. Each subfilter has $\sigma_c = 2°$.

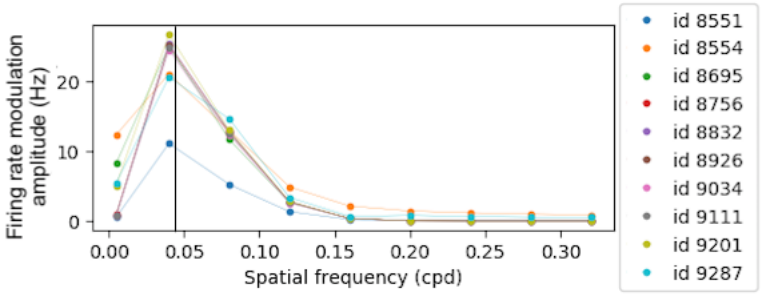

Figure B.13: Spatial frequency tuning curves of ten randomly selected sONtOFF neurons in the LGN. The curves show results from model simulations, while the vertical black line shows the analytical approximation used to map the preferred spatial frequencies of all LGN cells in histogram 9.10 section 9.2.3. This value (0.045 cpd) is the peak frequency obtained analytically when treating the neurons' receptive fields as a single, circular region. The sONtOFF neurons have receptive fields consisting of two DoG filters (equation (5.6)) of opposite polarity, separated by a few degrees. Each subfilter has $\sigma_c = 3°$.

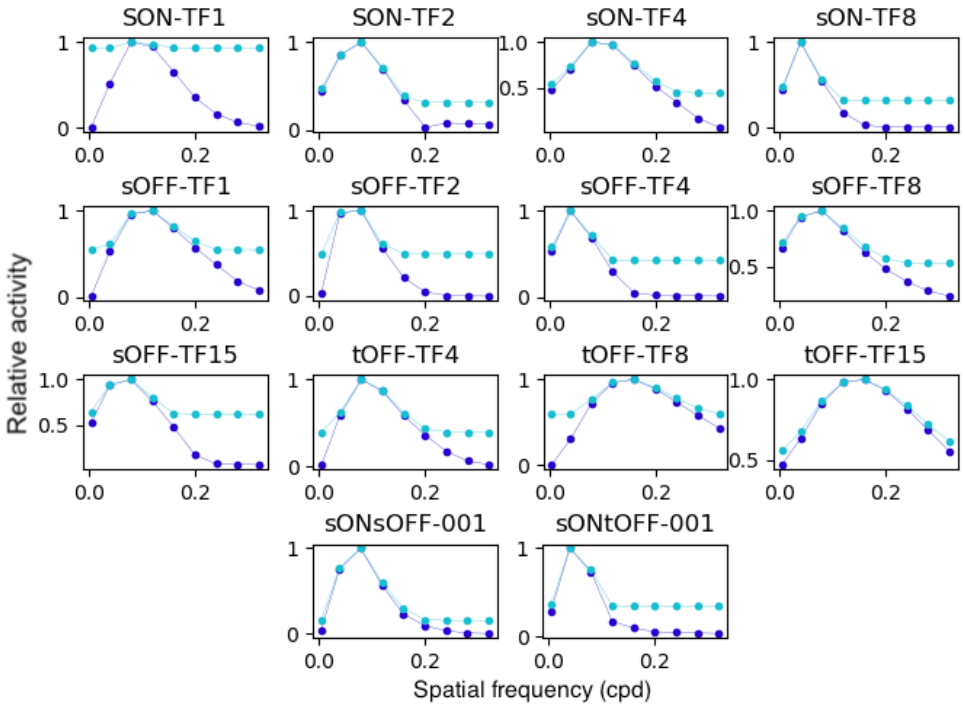

Figure B.14: Spatial frequency tuning of M1 LGN neurons, quantified by the normalised F0 (cyan) and F1 (dark blue) responses. The stimuli were full-field drifting gratings with a temporal frequency of 4 Hz.

V1

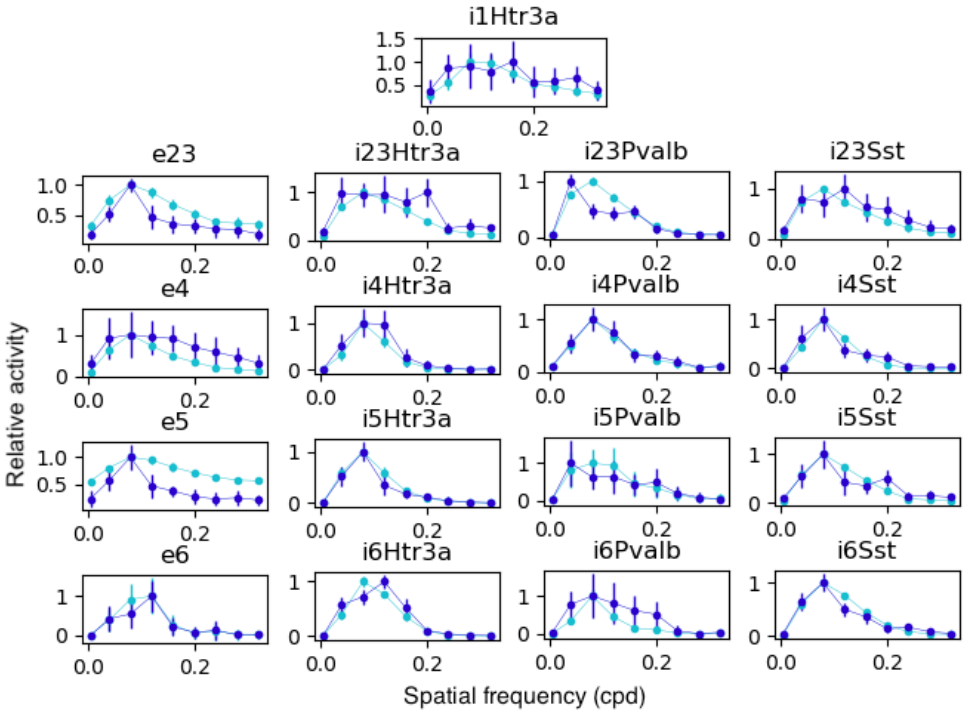

Figure B.15: Spatial frequency tuning of M1 V1 neurons, quantified by the normalised F0 (cyan) and F1 (dark blue) responses. The stimuli were full-field drifting gratings with a temporal frequency of 4 Hz. Curves show the mean over 20 trials of 3-second simulations, with error bars indicating the standard deviation.

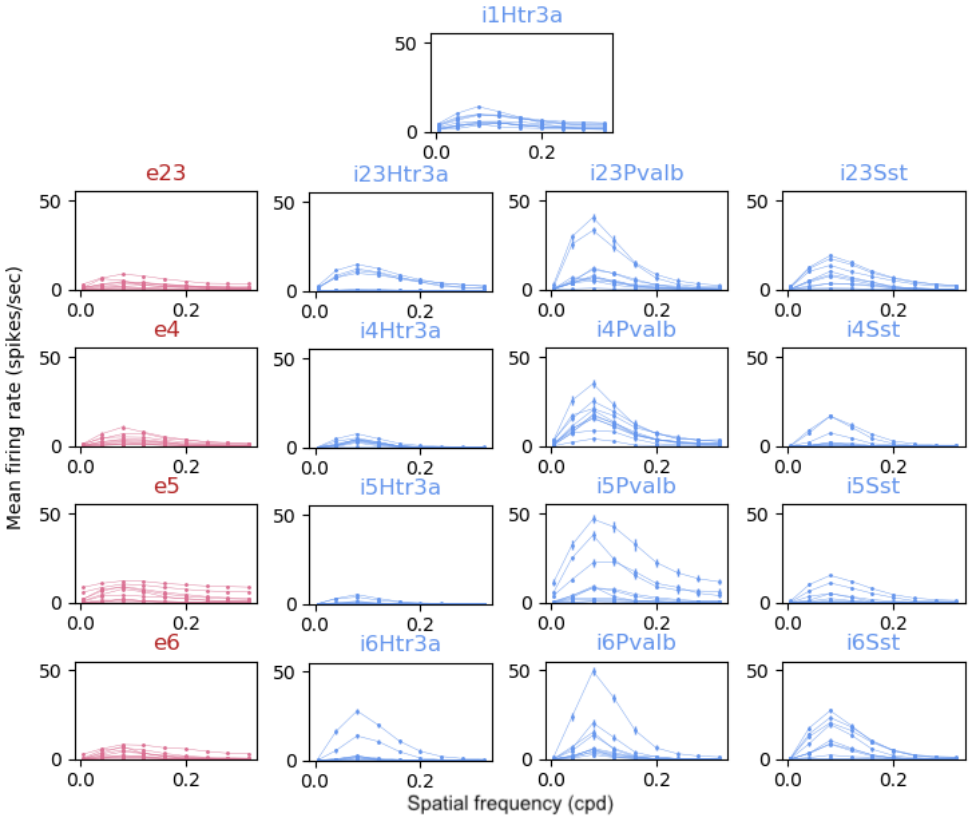

Figure B.16: Spatial frequency tuning of V1 neurons in the M1 Allen model. Each subplot shows 10 responsive neurons from the respective classes. The stimuli were full-field drifting gratings with a temporal frequency of 4 Hz. Curves show the mean over 20 trials of 3-second simulations, with error bars indicating the standard deviation.

B.2.2 Size tuning - flashing spots

LGN

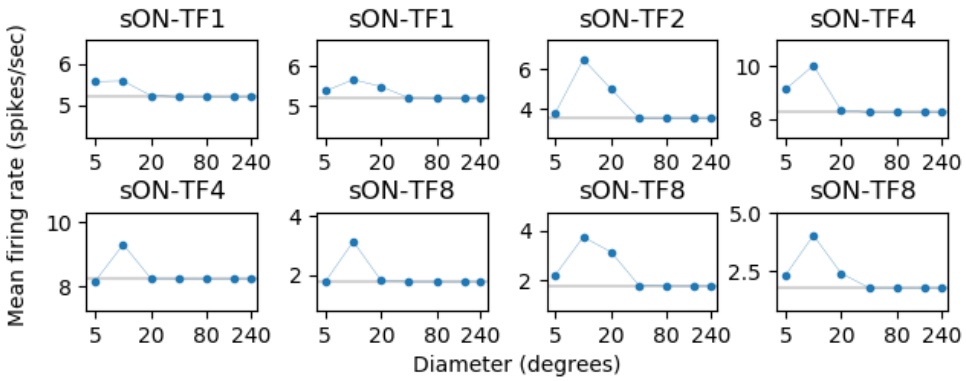

Figure B.17: Size tuning of eight ON-neurons in the M1 LGN, driven by 300 ms flashes of white spots. The tuning of one example neuron from each cell class is shown. The spontaneous rate of each neuron is indicated by the grey horizontal line.

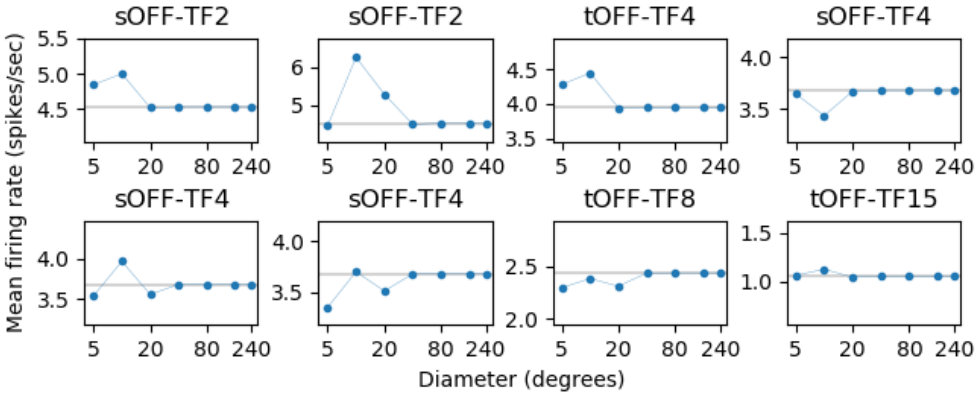

Figure B.18: Size tuning of eight OFF-neurons in the M1 LGN, driven by 300 ms flashes of white spots. The tuning of one example neuron from each cell class is shown. The spontaneous rate of each neuron is indicated by the grey horizontal line.

V1

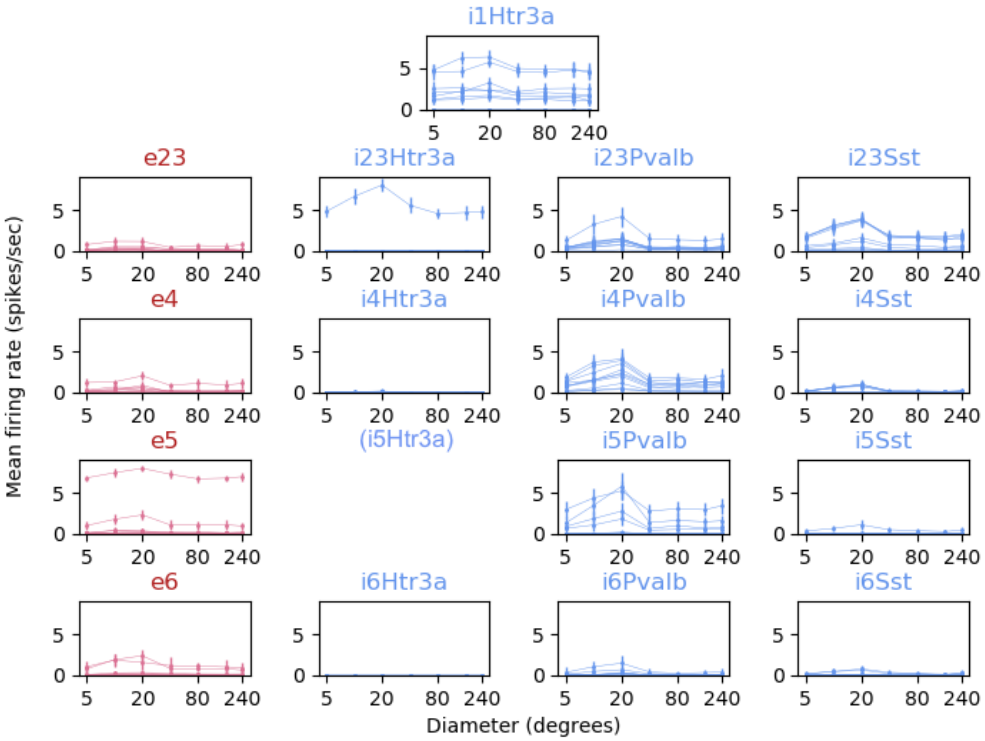

Figure B.19: Size tuning of V1 neurons from the M1 Allen model, using 75ms flashes of white spots as stimuli. Up to ten neurons from each population are included, depending on the number of cells within 70 µm of stimuli center. For the i5Htr3a class, there were no cells within this radius. The measurements show mean response over 20 trials of 3-second simulations, with error bars representing the standard deviation.

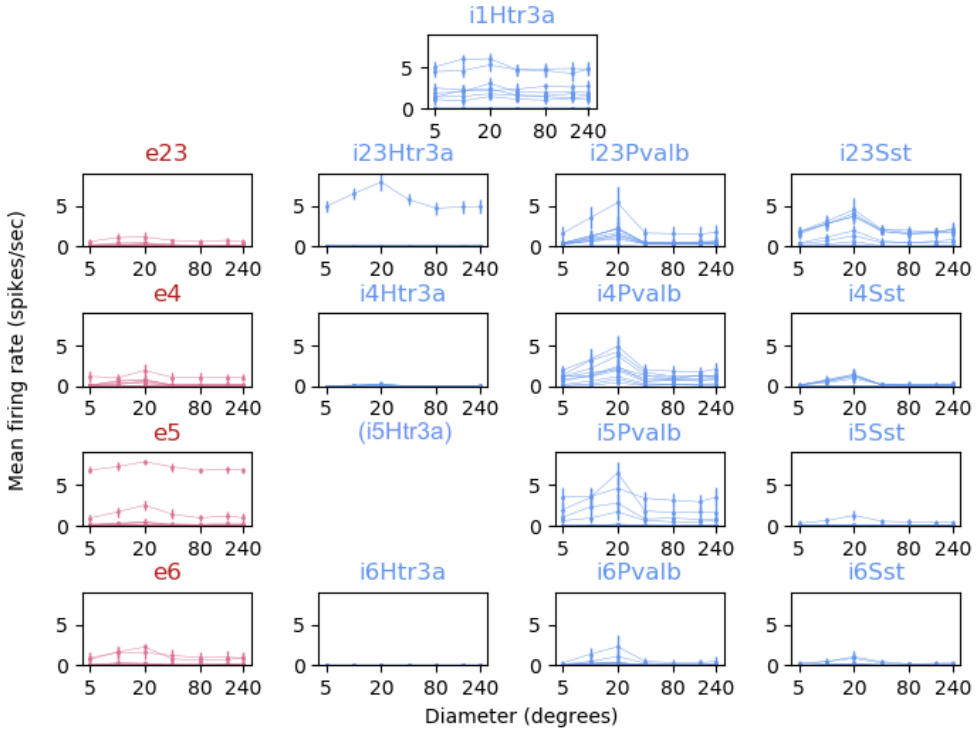

Figure B.20: Size tuning of V1 neurons from the M1 Allen model, using 300 ms flashes of white spots as stimuli. Up to ten neurons from each population are included, depending on the number of cells within 70 µm of stimuli center. For the i5Htr3a class, there were no cells within this radius. The measurements show mean response over 20 trials of 3-second simulations, with error bars representing the standard deviation.

B.2.3 Size tuning - patch gratings

LGN

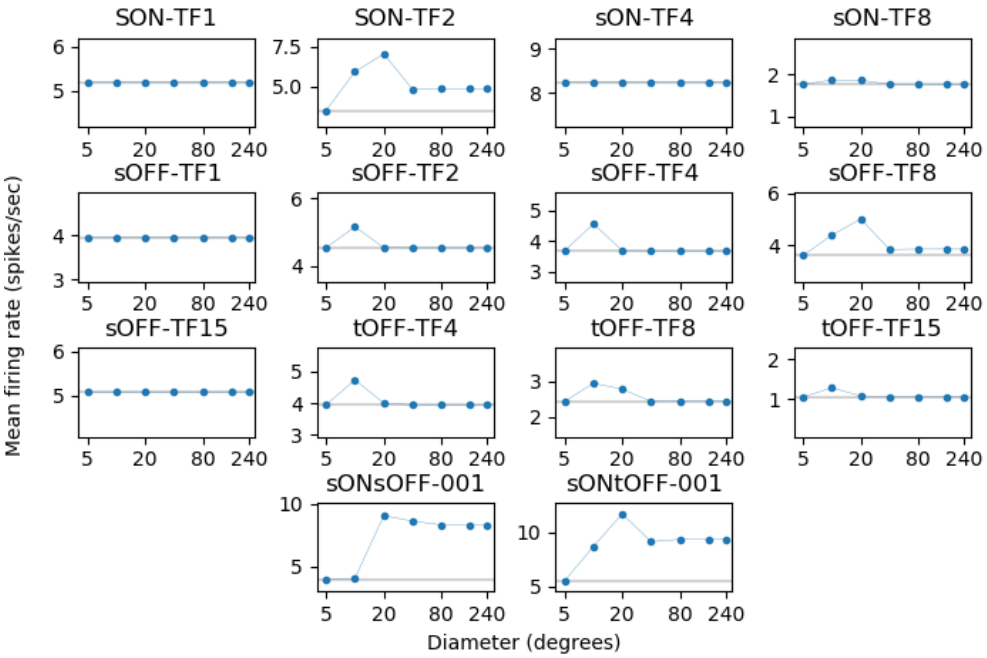

Figure B.21: Size tuning for one example neuron of each class in the M1 LGN. The stimuli were patch gratings with spatial and temporal frequencies of 0.02 cpd and 4 Hz. The spontaneous rate of the cell is marked by the grey line.

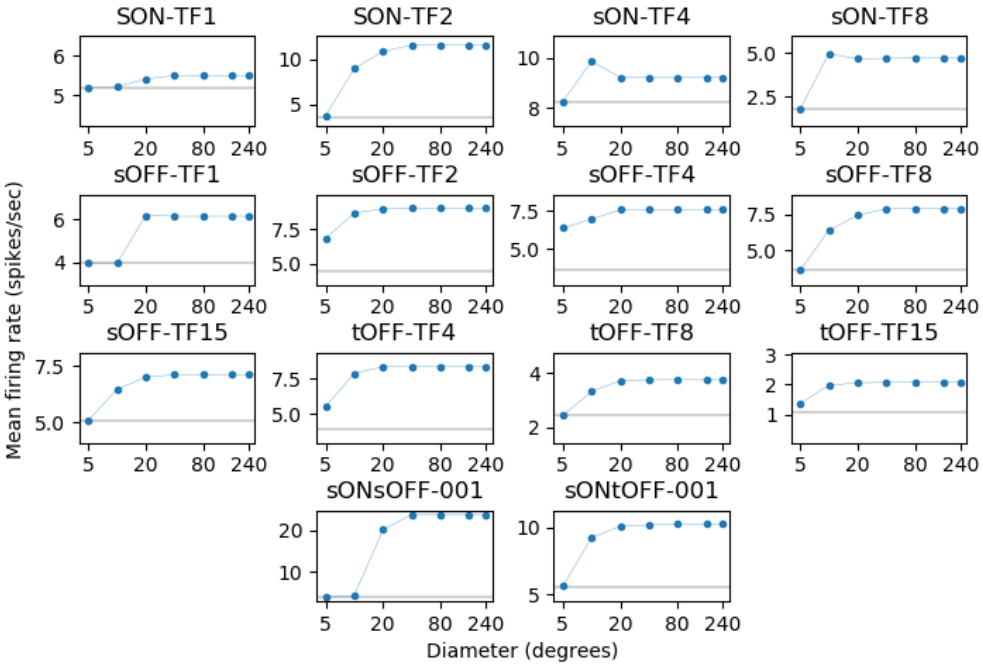

Figure B.22: Size tuning for one example neuron of each class in the M1 LGN. The stimuli were patch gratings with spatial and temporal frequencies of 0.08 cpd and 4 Hz. The spontaneous rate of the cell is marked by the grey line.

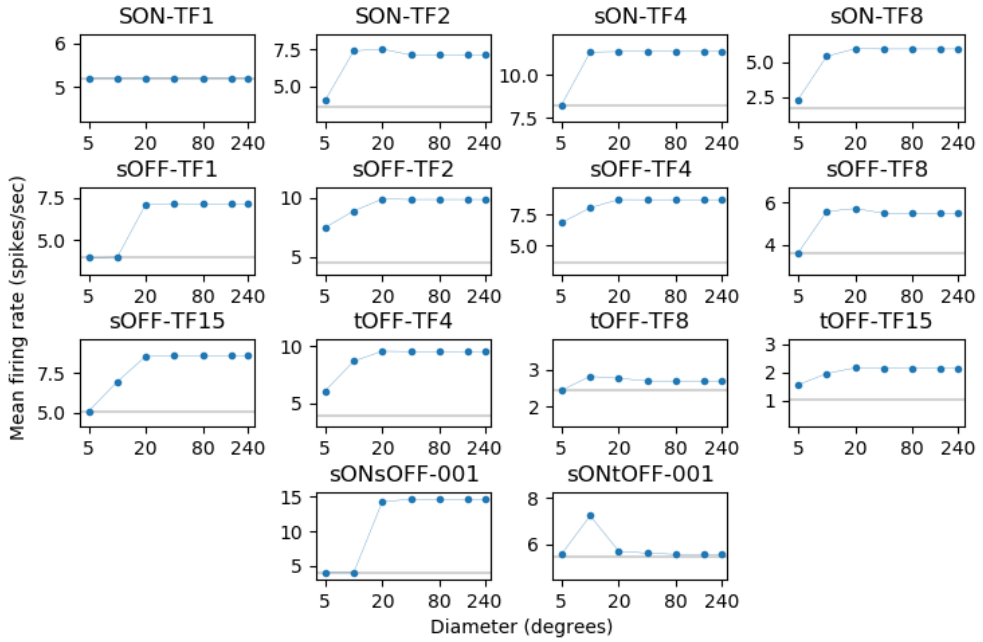

Figure B.23: Size tuning for one example neuron of each class in the M1 LGN. The stimuli were patch gratings with spatial and temporal frequencies of 0.12 cpd and 4 Hz. The spontaneous rate of the cell is marked by the grey line.

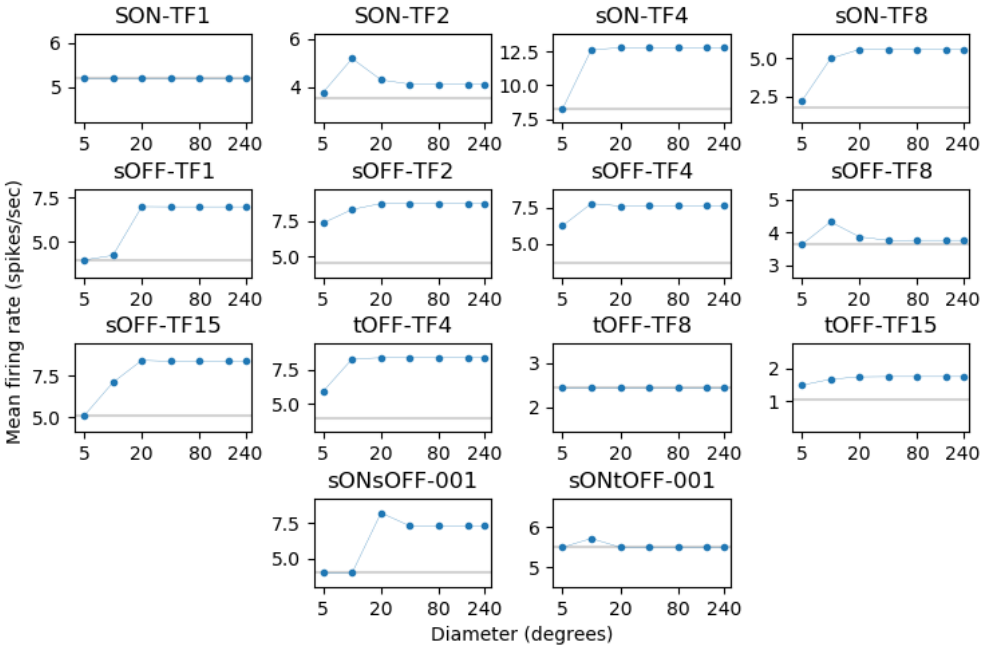

Figure B.24: Size tuning for one example neuron of each class in the M1 LGN. The stimuli were patch gratings with spatial and temporal frequencies of 0.16 cpd and 4 Hz. The spontaneous rate of the cell is marked by the grey line.

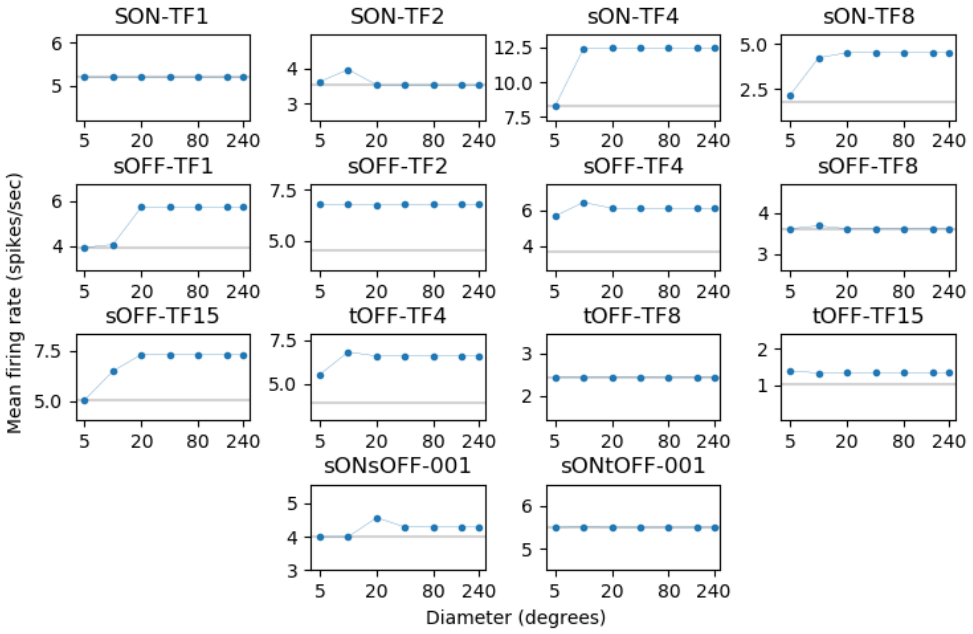

Figure B.25: Size tuning for one example neuron of each class in the M1 LGN. The stimuli were patch gratings with spatial and temporal frequencies of 0.2 cpd and 4 Hz. The spontaneous rate of the cell is marked by the grey line.

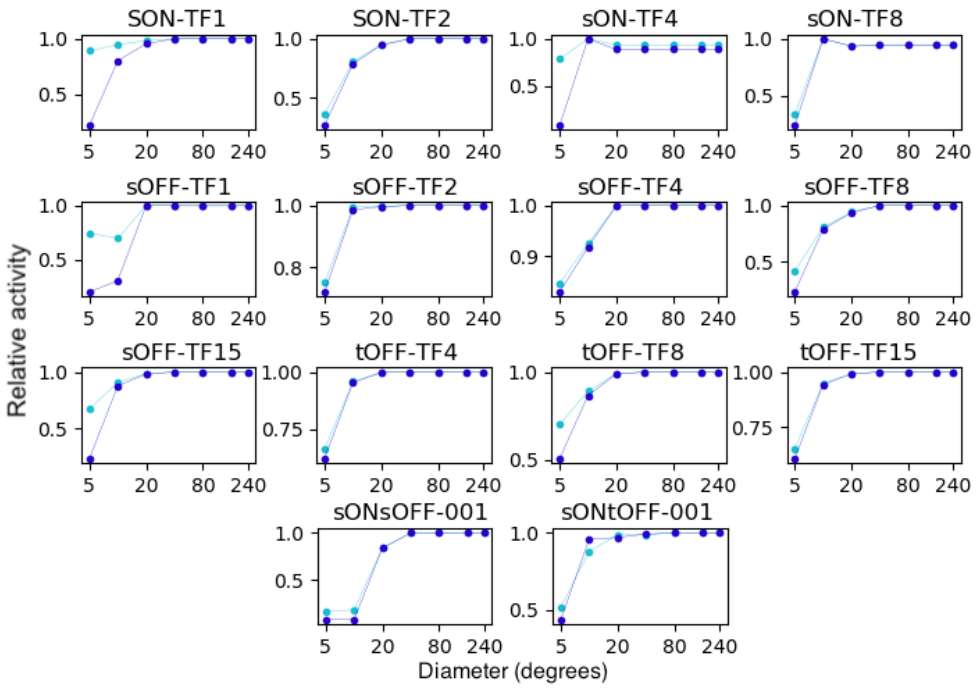

Figure B.26: Size tuning of M1 LGN neurons, quantified by the normalised F0 (cyan) and F1 (dark blue) responses. The stimuli were patch gratings with spatial and temporal frequencies of 0.08 cpd and 4 Hz.

V1

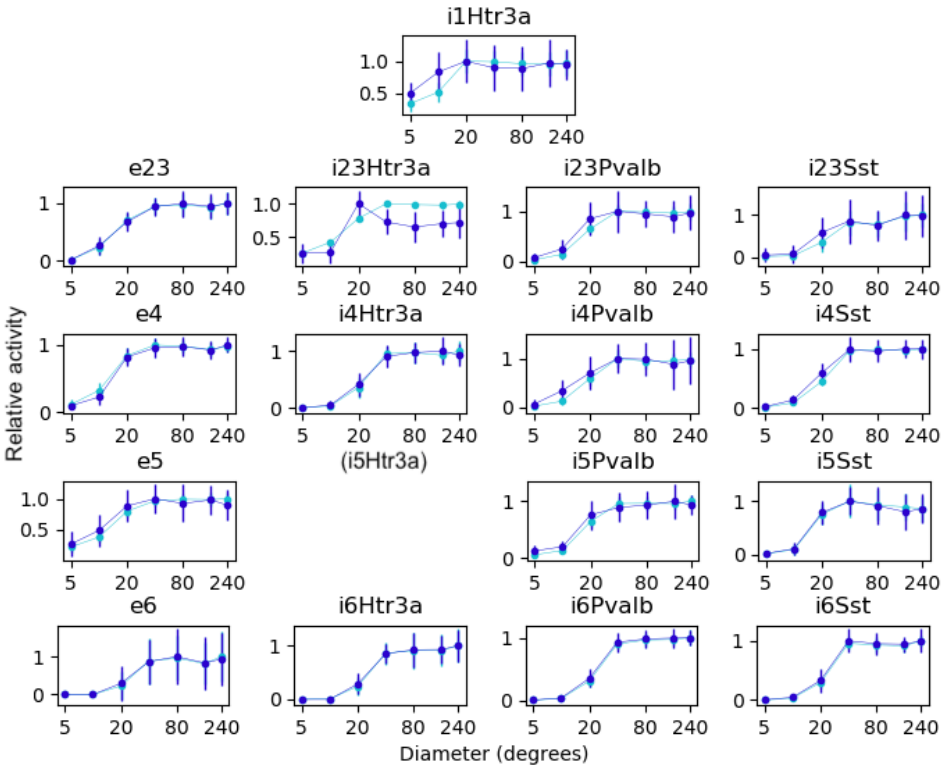

Figure B.27: Size tuning of M1 V1 neurons, quantified by the normalised F0 (cyan) and F1 (dark blue) responses. The stimuli were patch gratings with spatial and temporal frequencies of 0.08 cpd and 4 Hz. The neurons are selected from within 70µm of stimulus center. For the i5Htr3a class, there were no cells within this radius. The mean F1 and F0 responses are calculated over 20 trials for 3-second simulations. Error bars show the standard deviation of the mean.

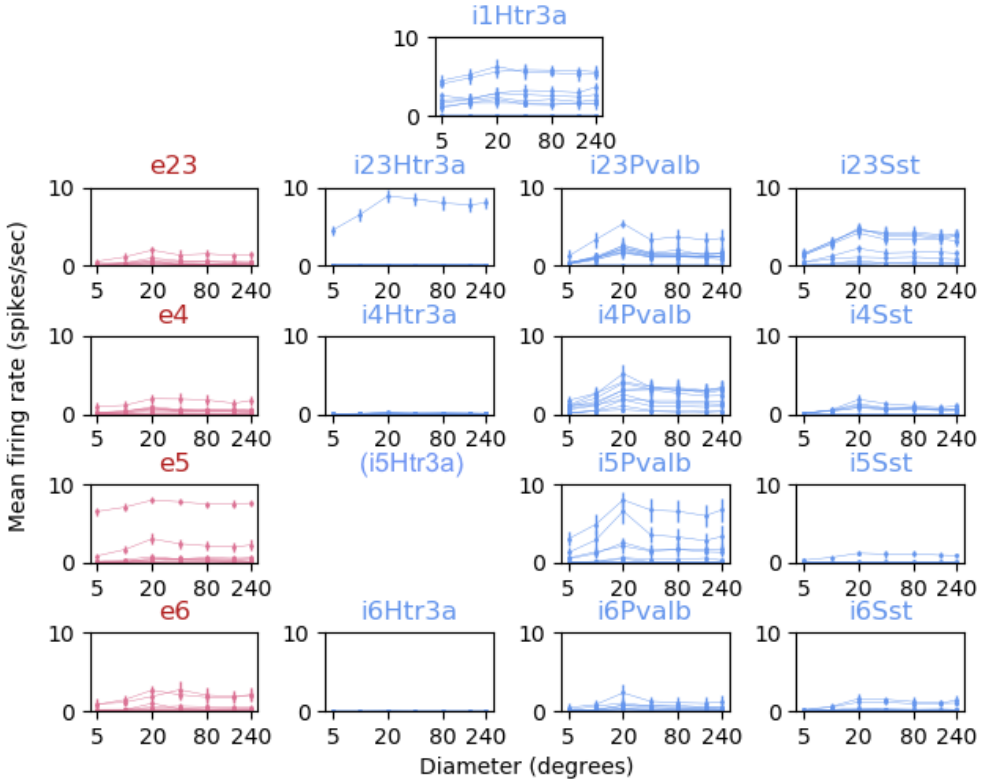

Figure B.28: Size tuning of V1 neurons from the Allen M1 model. The stimuli were patch gratings with spatial and temporal frequencies of 0.02 cpd and 4 Hz. Up to ten neurons from each population are included, depending on the number of cells within 70µm of stimulus center. For the i5Htr3a class, there were no cells within this radius. The mean is calculated over 20 trials for 3-second simulations. Error bars show the standard deviation of the mean.

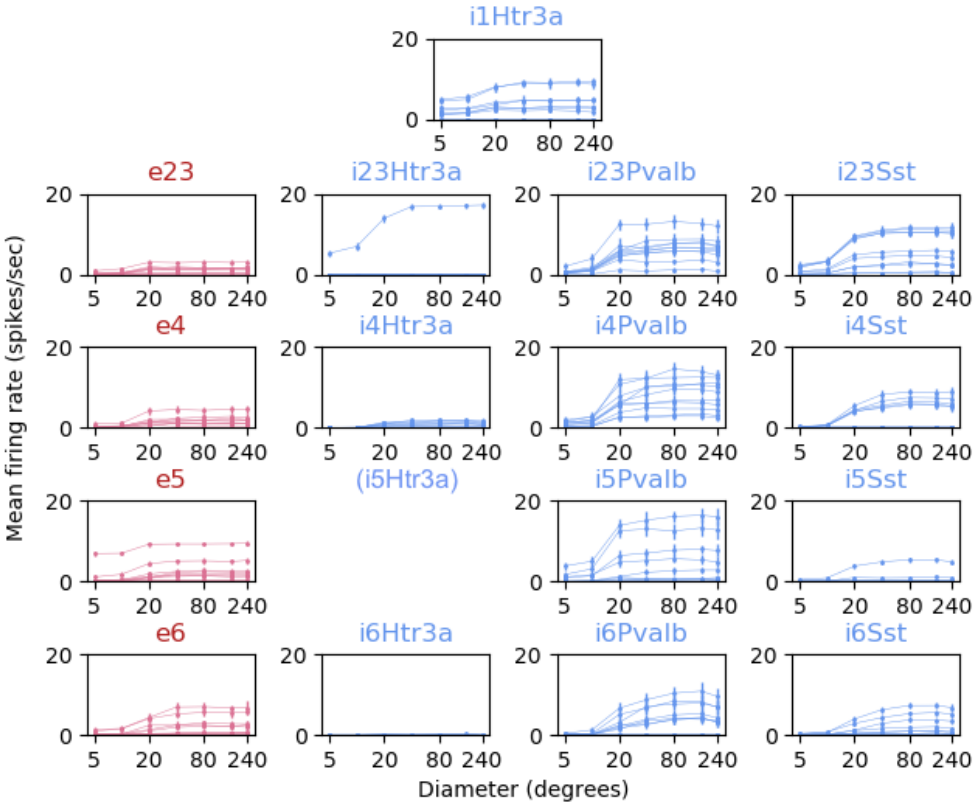

Figure B.29: Size tuning of V1 neurons from the Allen M1 model. The stimuli were patch gratings with spatial and temporal frequencies of 0.04 cpd and 4 Hz. Up to ten neurons from each population are included, depending on the number of cells within 70µm of stimulus center. For the i5Htr3a class, there were no cells within this radius. The mean is calculated over 20 trials for 3-second simulations. Error bars show the standard deviation of the mean.

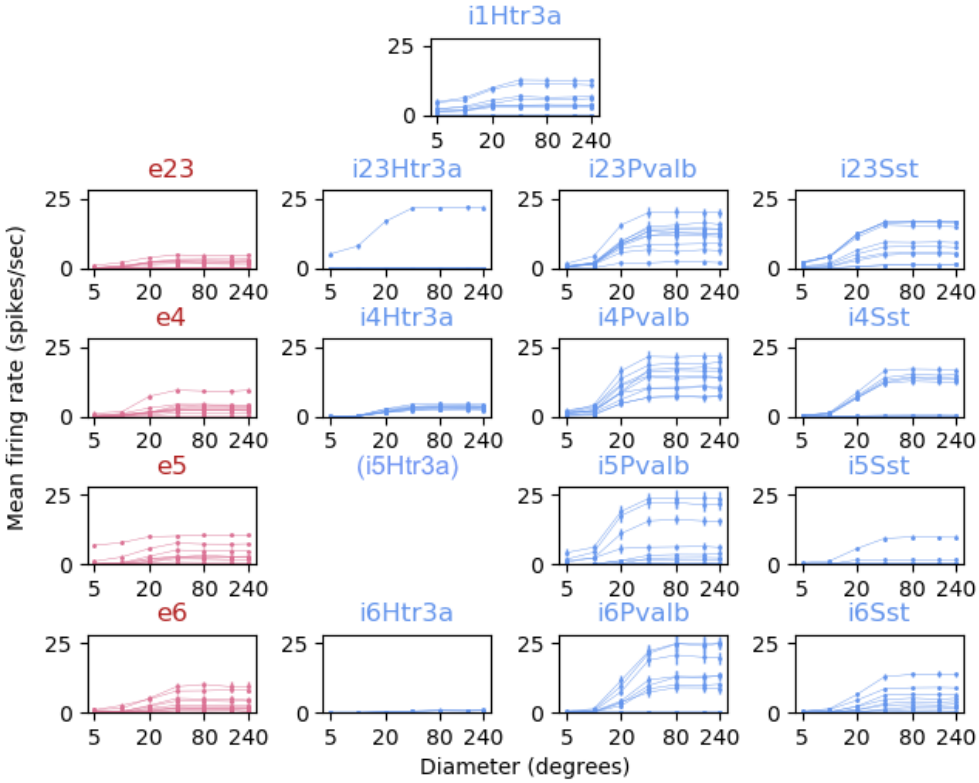

Figure B.30: Size tuning of V1 neurons from the Allen M1 model. The stimuli were patch gratings with spatial and temporal frequencies of 0.08 cpd and 4 Hz. Up to ten neurons from each population are included, depending on the number of cells within 70µm of stimulus center. For the i5Htr3a class, there were no cells within this radius. The mean is calculated over 20 trials for 3-second simulations. Error bars show the standard deviation of the mean.

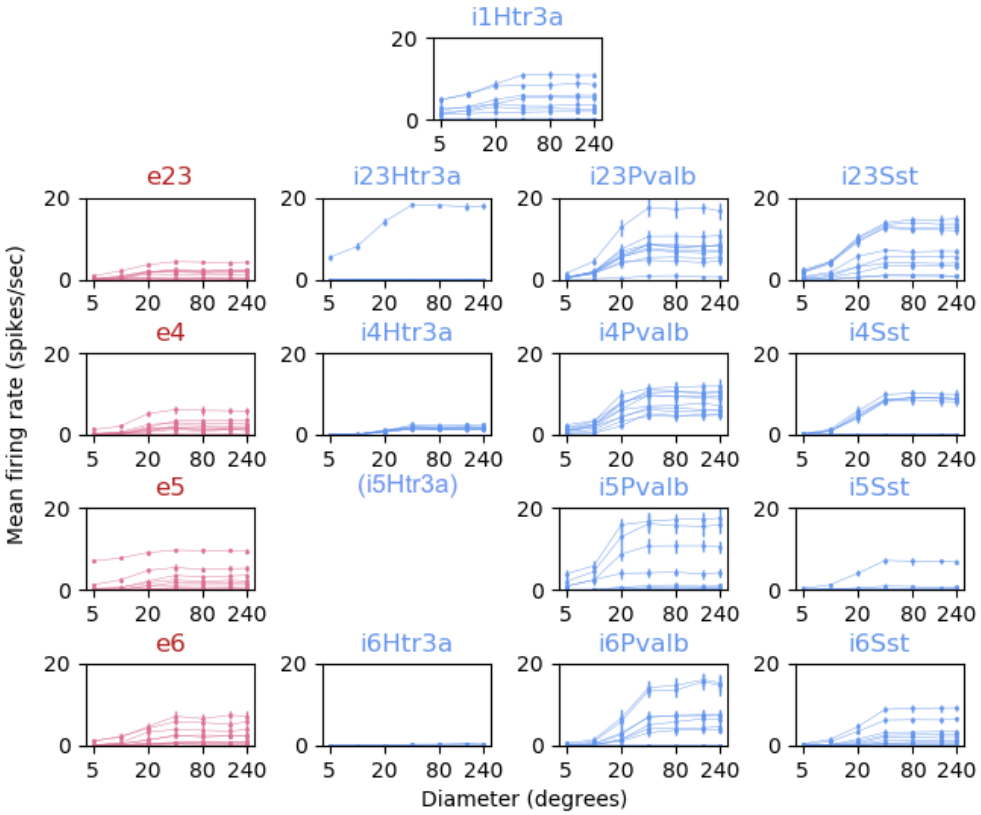

Figure B.31: Size tuning of V1 neurons from the Allen M1 model. The stimuli were patch gratings with spatial and temporal frequencies of 0.12 cpd and 4 Hz. Up to ten neurons from each population are included, depending on the number of cells within 70µm of stimulus center. For the i5Htr3a class, there were no cells within this radius. The mean is calculated over 20 trials for 3-second simulations. Error bars show the standard deviation of the mean.

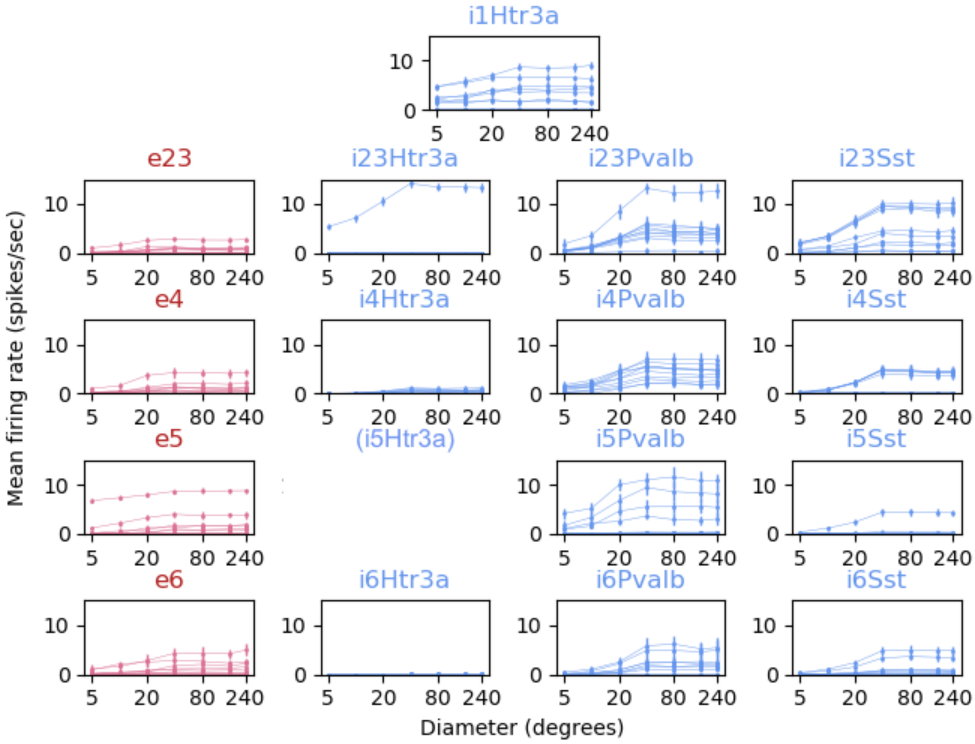

Figure B.32: Size tuning of V1 neurons from the Allen M1 model. The stimuli were patch gratings with spatial and temporal frequencies of 0.16 cpd and 4 Hz. Up to ten neurons from each population are included, depending on the number of cells within 70µm of stimulus center. For the i5Htr3a class, there were no cells within this radius. The mean is calculated over 20 trials for 3-second simulations. Error bars show the standard deviation of the mean.

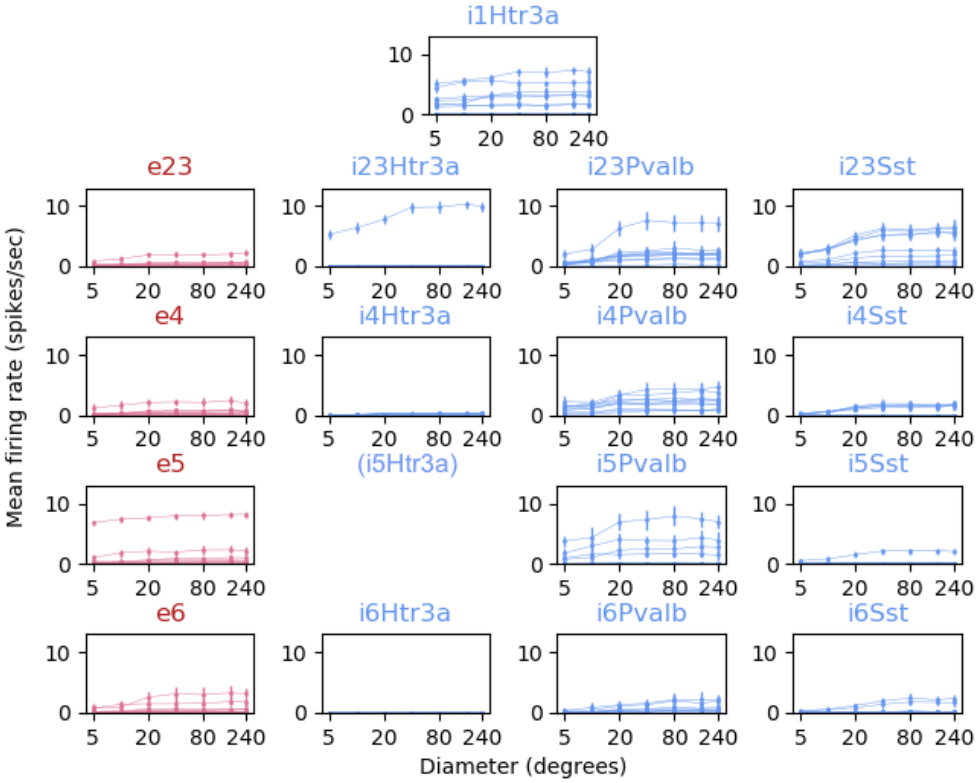

Figure B.33: Size tuning of V1 neurons from the Allen M1 model. The stimuli were patch gratings with spatial and temporal frequencies of 0.2 cpd and 4 Hz. Up to ten neurons from each population are included, depending on the number of cells within 70µm of stimulus center. For the i5Htr3a class, there were no cells within this radius. The mean is calculated over 20 trials for 3-second simulations. Error bars show the standard deviation of the mean.

B.3 DoG Filters with Various Parameters

Size tuning for patch grating in V1

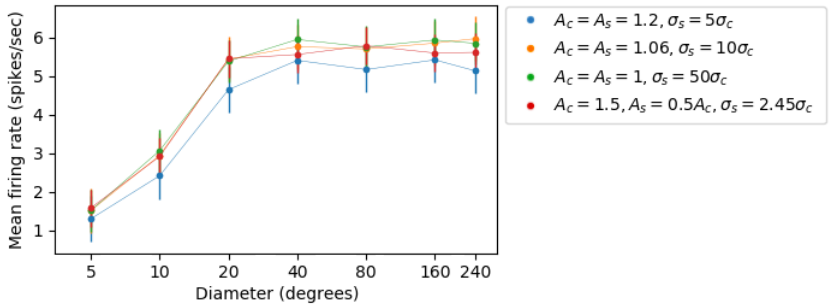

Figure B.34: Size tuning of a V1 neuron (id = 25515, type = e5). The labels show the parameters used for the DoG filters (equation (5.6)) in the LGN, which is used to drive the V1 model. Patches of drifting grating with spatial and temporal frequencies of 0.04 cpd and 4 Hz were used as stimuli. The mean is calculated over 20 trials for 3-second simulations. Error bars show the standard deviation of the mean.

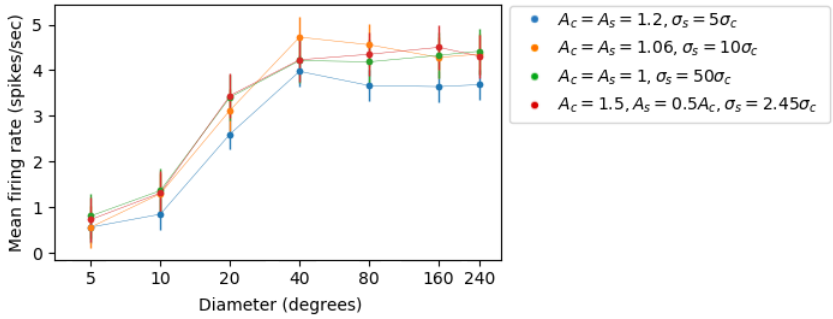

Figure B.35: Size tuning of a V1 neuron (id = 99626, type = e4). The labels show the parameters used for the DoG filters (equation (5.6)) in the LGN, which is used to drive the V1 model. Patches of drifting grating with spatial and temporal frequencies of 0.04 cpd and 4 Hz were used as stimuli. The mean is calculated over 20 trials for 3-second simulations. Error bars show the standard deviation of the mean.

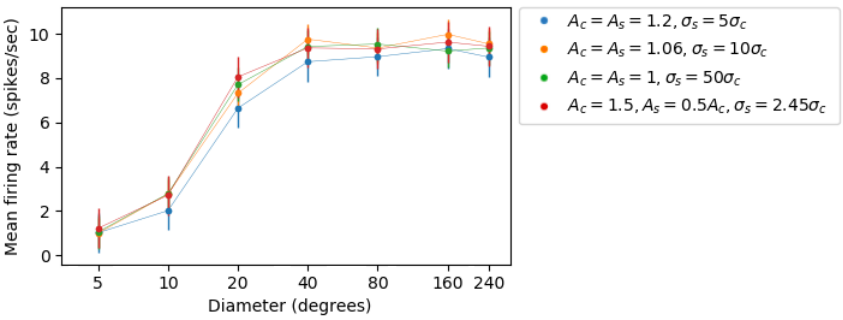

Figure B.36: Size tuning of a V1 neuron (id = 109331, type = i23Htr3a). The labels show the parameters used for the DoG filters (equation (5.6)) in the LGN, which is used to drive the V1 model. Patches of drifting grating with spatial and temporal frequencies of 0.04 cpd and 4 Hz were used as stimuli. The mean is calculated over 20 trials for 3-second simulations. Error bars show the standard deviation of the mean.

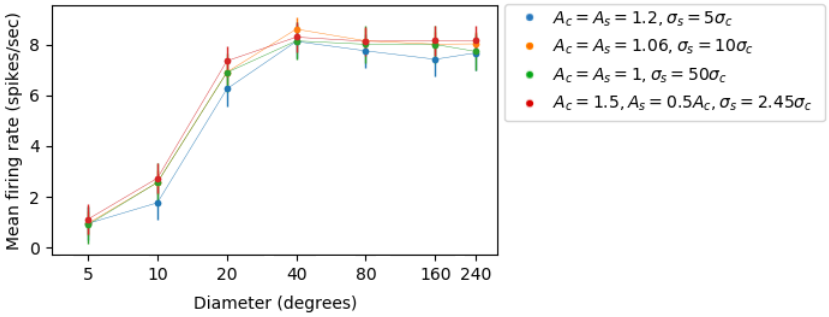

Figure B.37: Size tuning of a V1 neuron (id = 109362, type = i23Htr3a). The labels show the parameters used for the DoG filters (equation (5.6)) in the LGN, which is used to drive the V1 model. Patches of drifting grating with spatial and temporal frequencies of 0.04 cpd and 4 Hz were used as stimuli. The mean is calculated over 20 trials for 3-second simulations. Error bars show the standard deviation of the mean.

B.4 V1 Without Recurrent Connection

Size tuning with patch gratings

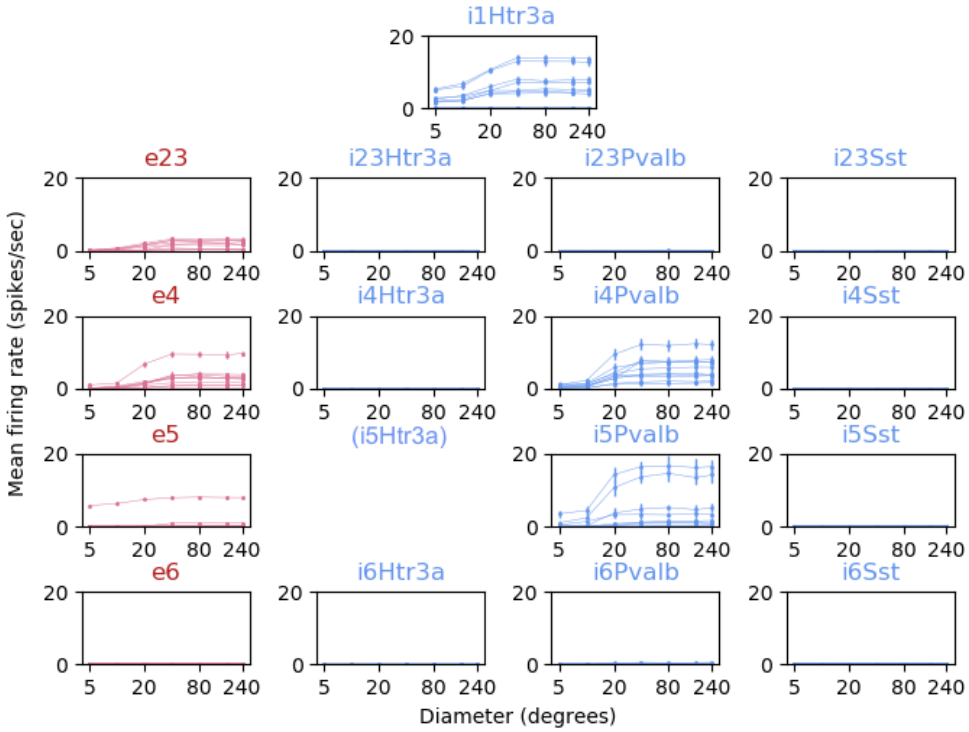

Figure B.38: Size tuning for cells in the $M1_{norec}$ V1 model. No response is generated by cells in the Sst and Htr3a classes in LII/III - LVI. These are the classes which do not receive LGN input. The i23Pvalb, i6Pvalb and E6 cells generate either weak or no activity. Most activity is seen for the i1Htr3a, i4Pvalb, i5Pvalb, and e4 cells. The tuning is measured in response to patch gratings with spatial and temporal frequencies of 0.08 cpd and 4 Hz. Up to ten neurons from each population are included, depending on the number of cells within 70µm of stimulus center. For the i5Htr3a class, there were no cells within this radius. The mean is calculated over 20 trials for 3-second simulations. Error bars show the standard deviation of the mean.

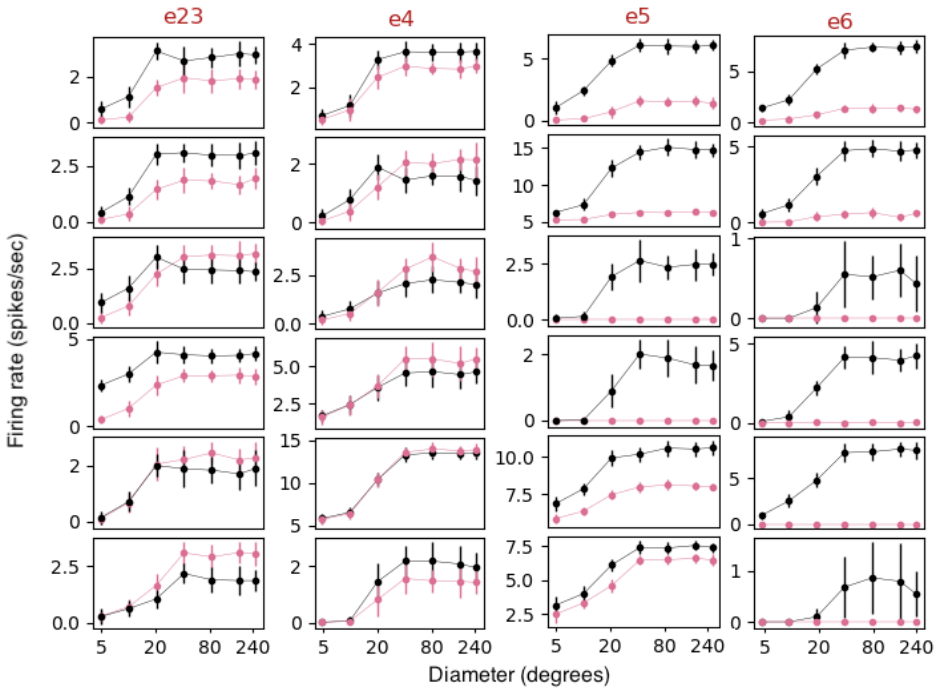

Figure B.39: Comparison of size tuning for two versions of the M1 V1 model: with (black) and without (red) recurrent connections. The tuning is measured in response to patch gratings with spatial and temporal frequencies of 0.08 cpd and 4 Hz. Six excitatory neurons is included from each layer, from LII/III to LVI. The mean is calculated over 20 trials for 3-second simulations. Error bars show the standard deviation of the mean.

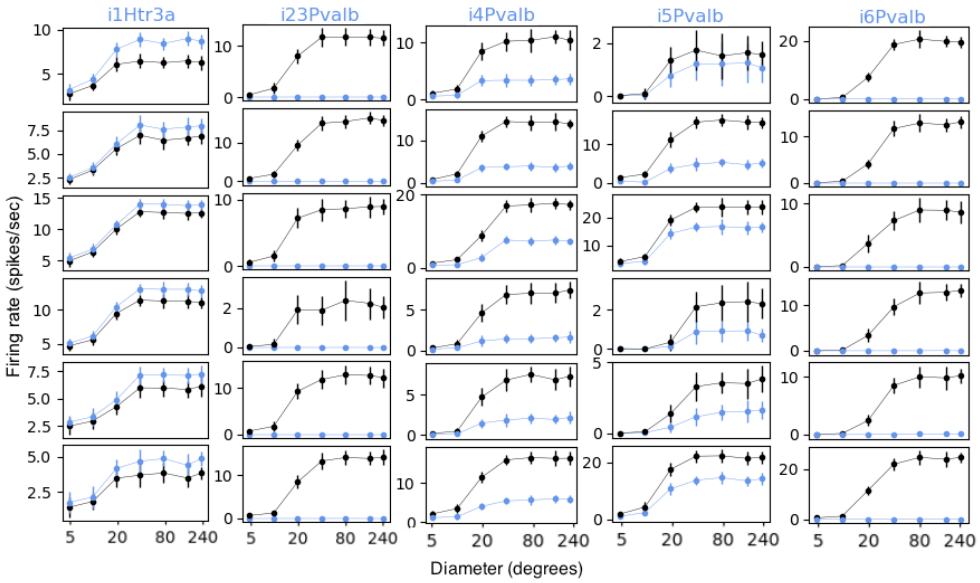

Figure B.40: Comparison of size tuning for two versions of the M1 V1 model: with (black) and without (blue) recurrent connections. The tuning is measured in response to patch gratings with spatial and temporal frequencies of 0.08 cpd and 4 Hz. Six inhibitory neurons is included from each layer, from LI to LVI. The mean is calculated over 20 trials for 3-second simulations. Error bars show the standard deviation of the mean.

B.5 Modelling nonlinear suppression

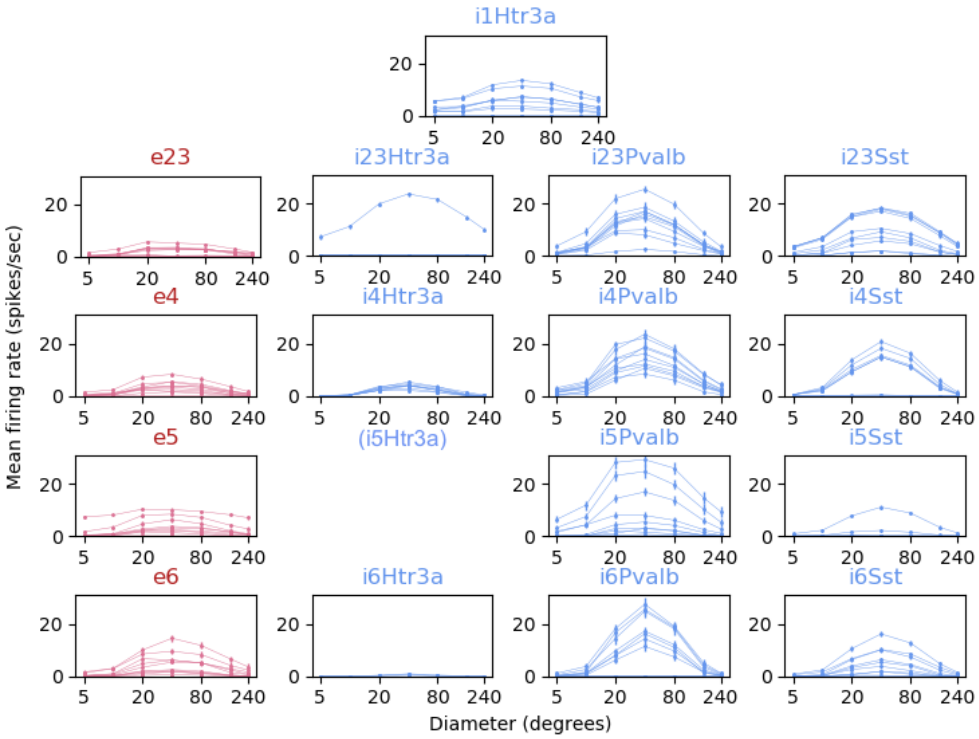

Figure B.41: Size tuning of V1 neurons for the Allen model with normalised LGN responses. The visual stimuli were increasing sizes of patch gratings with spatial and temporal frequencies of 0.04 cpd and 4 Hz. Up to ten neurons from each population are included, depending on the number of cells within 70µm of stimulus center. For the i5Htr3a class, there were no cells within this radius. The mean is calculated over 20 trials for 3-second simulations. Error bars show the standard deviation of the mean.

B.6 Temporal frequency tuning

LGN

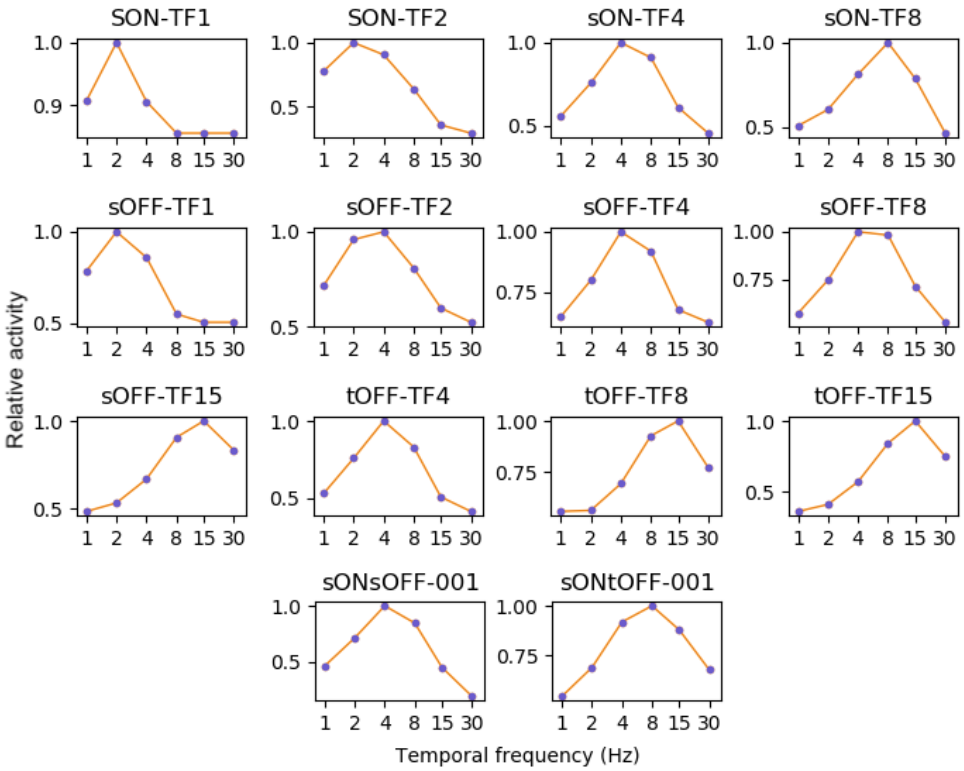

Figure B.42: Temporal frequency tuning of LGN neurons. The normalised tuning curves of the original (orange) and M1 (purple) model overlap. Full-field drifting gratings with spatial frequencies of 0.04 cpd (original) and 0.08 cpd (M1) are used as stimuli.

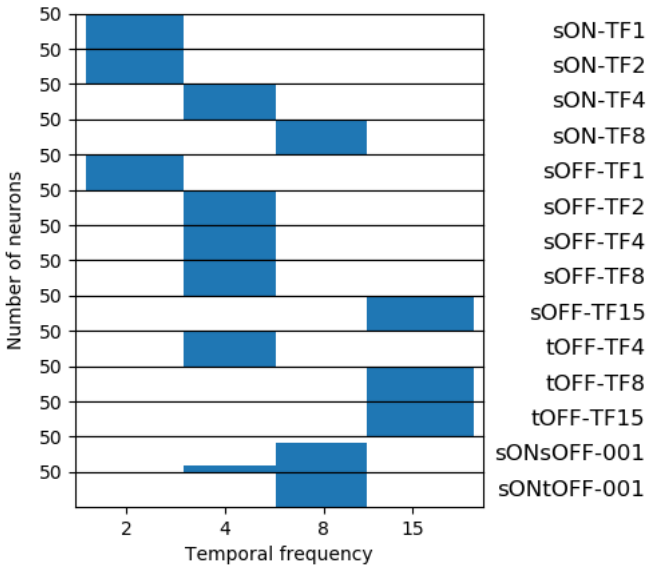

Figure B.43: Peak temporal frequencies LGN neurons, measured by their F1 responses. The stimulus is a full-field drifting grating, with spatial frequency of 0.08 cpd and temporal frequencies 1, 2, 4, 8, 15 and 30 Hz. 50 cells were analysed for each class. Not included in the figure are frequencies 1 and 30 Hz, to which none of the of the neurons showed a preference.

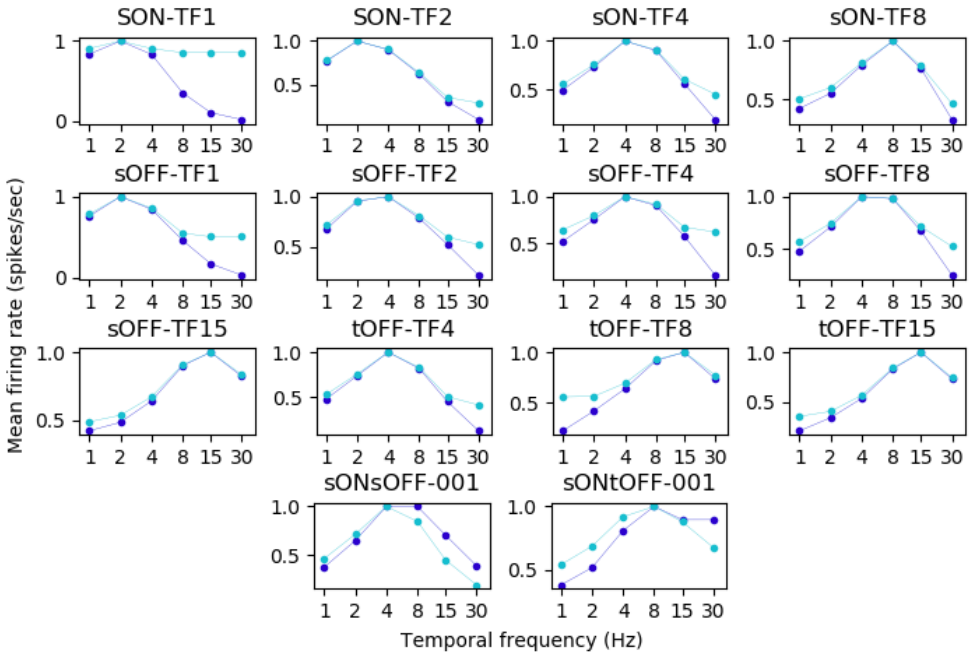

Figure B.44: Temporal frequency tuning of M1 LGN neurons, quantified by the normalised F0 (cyan) and F1 (dark blue) responses. The stimuli is a full-field drifting grating with spatial frequency of 0.08 cpd.

V1

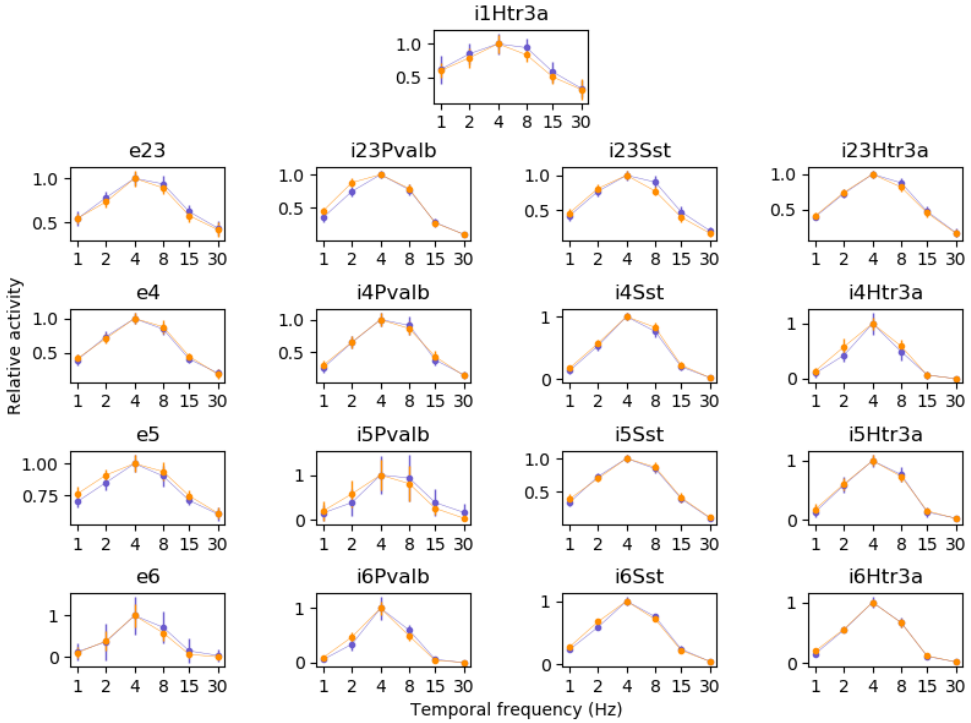

Figure B.45: Temporal frequency tuning of V1 neurons. The normalised tuning curves of the original (orange) and M1 (purple) model overlap. Full-field drifting gratings with spatial frequencies of 0.04 cpd (original) and 0.08 cpd (M1) are used as stimuli. Curves show the mean over 20 trials of 3-second simulations, with error bars indicating the standard deviation.

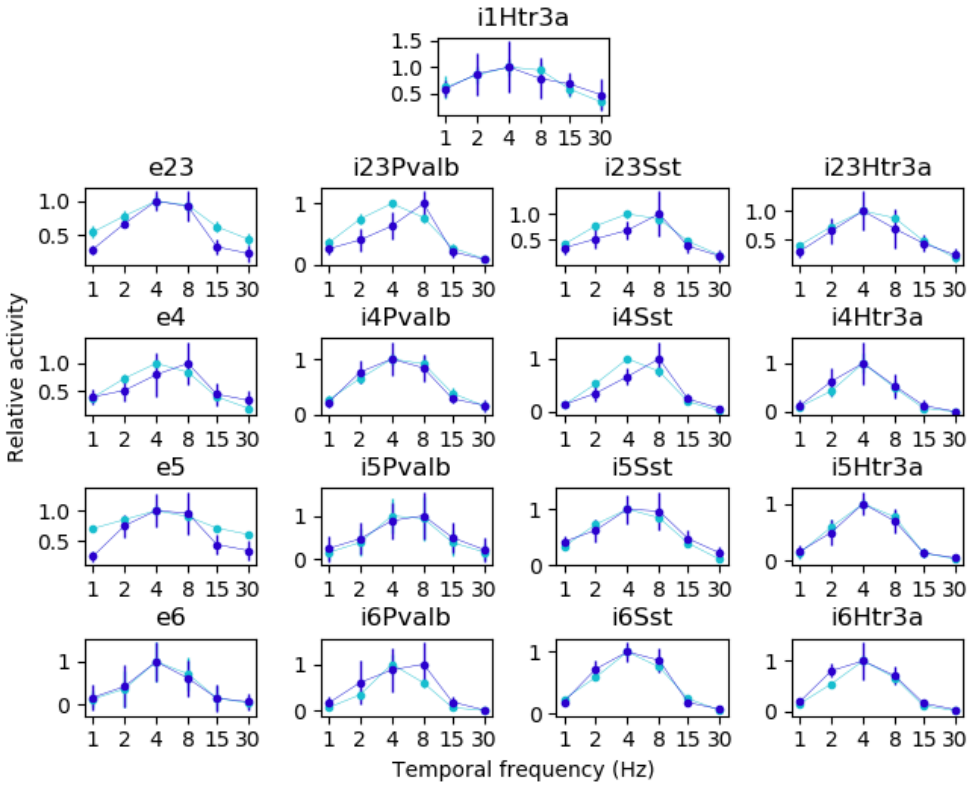

Figure B.46: Temporal frequency tuning of M1 V1 neurons, quantified by the normalised F0 (cyan) and F1 (dark blue) responses. The stimuli used are full-field drifting gratings with a spatial frequency of 0.08 cpd. Curves show the mean over 20 trials of 3-second simulations, with error bars indicating the standard deviation.

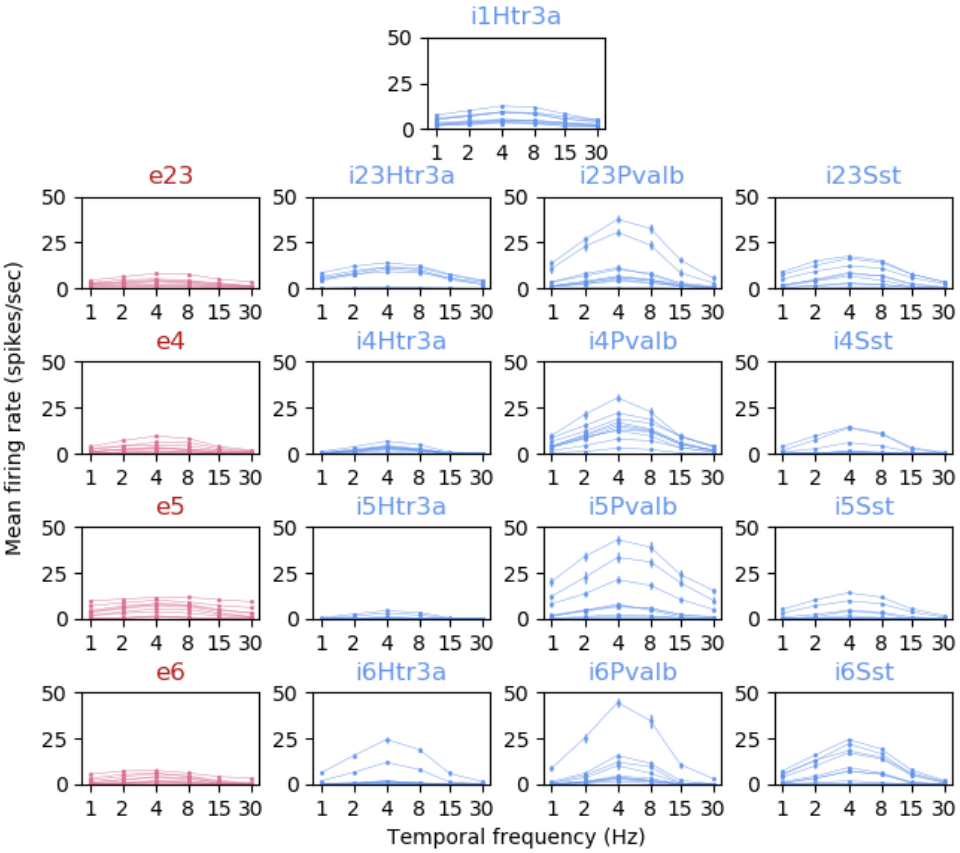

Figure B.47: Temporal frequency tuning of V1 neurons in the M1 Allen model. Each subplot shows 10 responsive neurons from the respective classes. The stimuli used are full-field drifting gratings with a spatial frequency of 0.08 cpd. Curves show the mean over 20 trials of 3-second simulations, with error bars indicating the standard deviation.

Appendix C

Code

The main code written for this thesis is available through GitHub at `https://github.com/lenamyk/Exploring_the_Allen_model`. The repository provides code used for analysis, creating visual input, and for creating the M1 and M2 models based on the original LGN. The modified files are included, along with an overview of the modifications made to each script. The original Allen model is available from: `https://www.dropbox.com/sh/xb7xasih3d8027u/AAAbKXe0Zmk86o3_y1iPVPCLa?dl=0`.

C.1 Example Code

Below, three scripts showing the main calculations for this thesis is demonstrated. The first presents how the preferred spatial frequencies of LGN cells were calculated from their receptive field parameters. The second and third demonstrate how firing rates and tuning curves were obtained from simulation outputs.

Calculate preferred spatial frequencies of LGN cells

Receptive field sizes are first loaded in for each LGN cell from a list of the "spatial size" parameter for each LGN cell. Here the list is contained in `spatialfilter.csv`, a subset of the file `lgn_full_col_cells_3.csv`[1]. The preferred spatial frequencies are then computed for all LGN cells, based on the analytical formula for spatial frequency tuning of a DoG filter (equation (5.9)).

```
#!/usr/bin/env python2
# -*- coding: utf-8 -*-

"""
Parameters
----------
nu : spatial frequency
k: wavevector
Ac: amplitude of center Gaussian
As: amplitude of surround Gaussian
spatial_sizes: receptive field sizes
               defined by the Allen model in
```

[1] See [112] for link to `lgn_full_col_cells_3.csv`.

```
                "LGN/LGN/lgn_full_col_cells_3.csv"
sigma_c: width of center Gaussian
         defined in the Allen model as the spatial size divided by 3
sigma_s: width of surround Gaussian
sc: scaling factor for determining sigma_s
"""

import numpy as np
import matplotlib.pyplot as plt
import math
import pandas as pd


Ac = 1.73
As = 1*Ac
sc = 2.45

spatial_sizes = pd.read_csv('spatialfilter.csv', delimiter='\t')
sigma_c = spatial_sizes.values/3
sigma_s = sc*sigma_c

nu = np.arange(0, 1, 0.001)
k = nu*2*np.pi
mean_f = np.zeros(len(peaks),len(k))  # mean frequency response
peaks = np.zeros(len(sigma_c))  # preferred frequency

# Calculate spatial frequency tuning for each cell:
for i in range(len(peaks)):
    ind = 0

    for j in k:
        mean_f[i,ind] = (Ac*math.exp((-j**2 *sigma_c[i]**2)/2)
                         - As*math.exp((-j**2 *(sigma_s[i])**2)/2))
        ind = ind + 1

    peaks_ind = np.argmax(mean_f[i,:])
    peaks[i] = nu[peaks_ind]

# Plot preferred frequencies:
plt.hist(peaks,
         bins=250,
         weights=np.zeros_like(peaks) + 1/peaks.size,
         range=(0, 0.25))
plt.xlabel('Spatial frequency (cpd)')
plt.ylabel('Fraction of cells')

# Find median, minimum and maximum preferred frequency:
median = np.median(peaks)
minimum = np.amin(peaks)
maximum = np.amax(peaks)
```

Calculate firing rates of LGN cells

Firing rates from LGN simulations were analysed for multiple neurons. The script below shows how firing rates for an example neuron were calculated and plotted as size tuning curves. The same methods were used for spatial and temporal frequency tuning. Scripts for the three cases are available on GitHub [2].

```
#!/usr/bin/env python3
# -*- coding: utf-8 -*-
"""
Parameters
----------
neuron_id: neuron to be analysed
simlength: length of simulation (ms)
tsteps_per_sec: temporal resolution (time steps per seconds)
tf: temporal frequency (Hz)
sf: spatial frequency (cpd)
tuning_array: set of stimuli diameters for the tuning curve
normalise_rates: set to 'True' if rates should be normalised
stim_type: type of stimuli, set to either 'flash' or 'grating'
grat_interval: measurement window for grating stimuli
flash_interval: measurement window for flashing stimuli
"""


import h5py
import numpy as np
import matplotlib.pyplot as plt


neuron_id = 8018
simlength = 3000
tsteps_per_sec = 1000
tf = 4
sf = 0.04
tuning_array = np.array([5,10,20,40,80,160,240])
normalise_rates = False
stim_type = 'grating'
grat_interval = 1000
flash_interval = 1200


# Firing rates (F0 and F1) for each stimuli size:
mean_F0 = np.zeros(len(tuning_array))
mean_F1 = np.zeros(len(tuning_array))


for i in range(len(tuning_array)):
```

[2] Code for calculating and visualising firing rates for size tuning, spatial frequency tuning and temporal frequency tuning: `https://github.com/lenamyk/Exploring_the_Allen_model/tree/master/code_and_analysis/analyse_simulation_output`

```
    # Read in firing rates from an example LGN output file:
    f = h5py.File('LGN/m1/%spix/3s_0.2gray_155d_4tf_0.04cpd.h5' \
                  % (tuning_array[i]), 'r')
    fr = f['firing_rates_Hz']
    t = f['time']
    fr = np.asarray(fr)
    t = np.asarray(t)
    time = np.arange(len(fr[1, :]))
    ids = np.arange(len(fr[:, 1]))


    # Set measurement window:
    if stim_type == 'grating':
        interval = grat_interval
    elif stim_type == 'flash':
        interval = flash_interval
    period_start = int(simlength/2)
    period_end = period_start + interval


    # Calculate F1:
    if stim_type == 'grating':
        signal = fr[neuron_id,period_start:period_end]
        ff = np.fft.fft(signal)
        ff_freq = np.fft.fftfreq(signal.size)*tsteps_per_sec
        amplitude_spectum = np.abs(ff)*2/(period_end - period_start)
        amp_ind = np.where(ff_freq == tf)
        mean_F1[i] = amplitude_spectum[amp_ind]


    # Calculate time average (or F0):
    mean_F0[i] = np.mean(fr[neuron_id, period_start:period_end])


# Normalise firing rates:
if normalise_rates and stim_type == 'grating':
    mean_F0 = mean_F0/np.amax(mean_F0)
    mean_F1 = mean_F1/np.amax(mean_F1)
    ylab = 'Relative activity'
else:
    ylab = 'Firing rate (spikes/sec)'


# Plot firing rates as a function of stimuli size:
if stim_type == 'grating':
    plt.plot(tuning_array, mean_F1, linestyle='-', marker='.',
             linewidth=0.2, c='mediumblue', label='F1')
    plt.plot(tuning_array, mean_F0, linestyle='-', marker='.',
             linewidth=0.2, c='tab:cyan', label='F0')
else:
    plt.plot(tuning_array, mean_F0, linestyle='-', marker='.',
```

```
            linewidth=0.2, c='tab:cyan', label='Mean firing rate ')
plt.xlabel('Diameter (degrees)')
plt.ylabel(ylab)
plt.figlegend()
```

Calculate firing rates of V1 cells

Spike trains from V1 simulations were analysed for multiple neurons in V1. The code below shows how firing rates were visualised and calculated, by using the size tuning of a single neuron as an example. The same methods were used for spatial and temporal frequency tuning. Scripts for the three cases are available on GitHub [3]. The spike trains in the code below are read in from a list of spike times for the selected neuron. The list was created by extracting the spike times from the V1 simulation output. A bash script which shows how spike trains were extracted is included in the same GitHub repository.

```
#!/usr/bin/env python2
# -*- coding: utf-8 -*-

"""
Parameters
----------
nrid: neuron to be analysed
simlength: length of simulation (ms)
tsteps_per_sec: temporal resolution (time steps per seconds)
trials: array of trial numbers
sf: spatial frequency (cpd)
tf: temporal frequency (Hz)
stim_type: type of stimuli, set to either 'flash' or 'grating'
dogparam: parameters of DoG filter
tuning_array: set of stimuli diameters for the tuning curve
normalise_rates: set to 'True' if rates should be normalised
"""


import numpy as np
import matplotlib.pyplot as plt
import statistics

nrid = '103609'
simlength = 3000
tsteps_per_sec = 1000
trials = np.arange(20)
sf = '0.08'
tf = 4
stim_type = 'grating'
dogparam = '1.73ac_1as_2.45sc'
```

[3] Code for calculating and visualising firing rates for size tuning, spatial frequency tuning and temporal frequency tuning: `https://github.com/lenamyk/Exploring_the_Allen_model/tree/master/code_and_analysis/analyse_simulation_output`

```
tuning_array = np.array([5,10,20,40,80,160,240])
normalise_rates = False


#Firing rates (F0 and F1) for each stimuli:
mean_F0 = np.zeros(len(tuning_array))
sd_F0 = np.zeros(len(tuning_array))
mean_F1 = np.zeros(len(tuning_array))
sd_F1 = np.zeros(len(tuning_array))


for j in range(len(tuning_array)):
    matrix_of_firing_rates = np.zeros((simlength,len(trials)))
    time_average_per_trial = np.zeros(len(trials))
    mean_amplitude_per_trial = np.zeros(len(trials))
    fig = plt.figure()


    for i in trials:
        # Read in spike output:
        filename = "m1/id%s/%dpix/3s_155d_4tf_0.08cpd_%s_tr%d.txt" \
        % (nrid,  tuning_array[j], dogparam, i)
        a = open(filename, 'r+')
        spike_times = a.readlines()
        spike_times = np.array(spike_times, dtype=float)
        a.close()


        #Create discrete firing rate:
        t = np.linspace(0, simlength + 1, simlength)
        firing_rate = np.zeros_like(t)
        entries = np.digitize(spike_times, t)
        firing_rate[entries] = 1*tsteps_per_sec
        matrix_of_firing_rates[:,i] = firing_rate


        #Calculate time average (or F0) per trial:
        time_average_per_trial[i] \
        = np.mean(matrix_of_firing_rates[200:, i])


        #Calculate F1 per trial:
        if stim_type == 'grating':
            start_measuring = 500
            signal = firing_rate[start_measuring:simlength]
            ff = np.fft.fft(signal)
            ff_freq = np.fft.fftfreq(signal.size)*tsteps_per_sec
            amp_spectr = np.abs(ff)*2/(simlength - start_measuring)
            ff_freq_pos = ff_freq[:int(len(ff_freq)/2)]
            amp_spectr_pos = amp_spectr[:int(len(amp_spectr)/2)]
            mean_amplitude_per_trial[i] = np.interp(tf,
```

```
                                                    ff_freq_pos,
                                                    amp_spectr_pos)


        # Plot firing rate for each trial:
        title = "trial %d" % (i + 1)
        ax = fig.add_subplot(len(trials) + 1, 1, i + 1)
        ax.set_title(title, loc='left', x=1.01, y=0)
        ax.plot(t, matrix_of_firing_rates[:,i], linewidth=0.5)
        ax.set_xticklabels([])


    #Calculate trial averages of firing rates:
    mean_spikes = np.mean(matrix_of_firing_rates, axis=1)
    mean_F0[j] = np.mean(mean_spikes)
    sd_F0[j] = statistics.stdev(time_average_per_trial)
    if stim_type == 'grating':
        mean_F1[j] = np.mean(mean_amplitude_per_trial)
        sd_F1[j] = statistics.stdev(mean_amplitude_per_trial)


    # Plot trial averaged firing rate:
    ax = fig.add_subplot(len(trials)+1, 1, len(trials)+1)
    ax.plot(t, mean_spikes,c='tab:orange', linewidth=0.5)
    ax.set_title('mean', loc='left', x=1.01, y=0)
    ax.set_title('%s degrees diameter' %(tuning_array[j]), y=21)
    ax.set_xlabel('Time (ms)')
    plt.subplots_adjust(hspace=0)
    plt.tight_layout()


#Normalise rates:
if normalise_rates and stim_type == 'grating':
    sd_F1 = sd_F1/np.amax(mean_F1)
    sd_F0 = sd_F0/np.amax(mean_F0)
    mean_F1 = mean_F1/np.amax(mean_F1)
    mean_F0 = mean_F0/np.amax(mean_F0)
    ylab = 'Relative activity'
else:
    ylab = 'Firing rate (spikes/sec)'


# Plot firing rates as a function of stimuli size:
fig = plt.figure()
if stim_type == 'grating':
    plt.errorbar(tuning_array, mean_F0, yerr=sd_F0, marker='.',
                 label='F0', c='tab:cyan')
    plt.errorbar(tuning_array, mean_F1, yerr=sd_F1, marker='.',
                 label='F1', c='mediumblue')
else:
    plt.errorbar(tuning_array, mean_F0, yerr=sd_F0, marker='.',
```

```
                label='Mean firing rate',c='tab:cyan')
plt.xlabel('Diameter (degrees)')
plt.ylabel(ylab)
plt.xscale('log')
plt.figlegend()
```