



Pattern formation in neural circuits by the
interaction of travelling waves with spike-timing
dependent plasticity

James E. M. Bennett

Thesis submitted for the degree of Doctor of Philosophy
Department of Physiology, Anatomy and Genetics
University of Oxford

Hilary 2014

Abstract

Spontaneous travelling waves of neuronal activity are a prominent feature throughout the developing brain and have been shown to be essential for achieving normal function, but the mechanism of their action on post-synaptic connections remains unknown. A well-known and widespread mechanism for altering synaptic strengths is spike-timing dependent plasticity (STDP), whereby the temporal relationship between the pre- and post-synaptic spikes determines whether a synapse is strengthened or weakened. Here, I answer the theoretical question of how these two phenomenon interact: what types of connectivity patterns can emerge when travelling waves drive a downstream area that implements STDP, and what are the critical features of the waves and the plasticity rules that shape these patterns? I then demonstrate how the theory can be applied to the development of the visual system, where retinal waves are hypothesised to play a role in the refinement of downstream connections. My major findings are as follows. (1) Mathematically, STDP translates the correlated activity of travelling waves into coherent patterns of synaptic connectivity; it maps the spatiotemporal structure in waves into a spatial pattern of synaptic strengths, building periodic structures into feedforward circuits. This is analogous to pattern formation in reaction diffusion systems. The theory reveals a role for the wave speed and time scale of the STDP rule in determining the spatial frequency of the connectivity pattern. (2) Simulations verify the theory and extend it from one-dimensional to two-dimensional cases, and from simplified linear wavefronts to more complex realistic and noisy wave patterns. (3) With appropriate constraints, these pattern formation abilities can be harnessed to explain a wide range of developmental phenomena, including how receptive fields (RFs) in the visual system are refined in size and topography and how simple-cell and direction selective RFs can develop. The theory is applied to the visual system here but generalises across different brain areas and STDP rules. The theory makes several predictions that are testable using existing experimental paradigms.

Contents

1	Introduction	7
2	General Methods	33
	Overview	34
	Spiking neuron simulation	34
	Output activity and synaptic plasticity	34
	Travelling waves in the input layer	37
	Computing numerical solutions	39
	Retinal wave analysis	40
	Converting between visual angle and retinal distance	43
3	Theoretical principles for spike-timing dependent plasticity of neural circuits driven by travelling wave activity patterns	47
	Background	47
	Methods	50
	Conceptual framework	50
	Travelling waves in the input layer	52
	Theoretical results	53
	Deriving synaptic strength dynamics	53
	Limitations to the induction of periodic patterns by waves and STDP . . .	58
	Simulation results	58
	Travelling waves induce spatial pattern formation in a spiking neuron simulation	58
	Pattern formation over a wide range of spatial and temporal scales . .	61
	Impact of burst duration on the separation of time scales	65
	Sensitivity of periodic connectivity patterns to noise	68
	Sensitivity of periodic connectivity patterns to multiple, non-isolated waves	71
	Discussion	75
	Appendix	79
4	Pattern formation as a mechanism to refine receptive fields	81
	Background	82
	Methods	85
	Travelling plane waves in the input layer	85
	Complex wave patterns	85

Receptive fields in the spiking neuron simulation	86
Results	87
Travelling waves are instructive for shaping receptive fields via spike-timing dependent plasticity	87
A quantitative relationship between patterning and refinement of receptive fields	92
Refinement of receptive fields in two spatial dimensions	96
Refinement of receptive fields by complex wave patterns	100
Discussion	104
5 Analysis of candidate STDP rules for retinocollicular refinement and their influence on the role played by retinal waves	111
Background	111
Methods	115
Retinal wave simulations	115
Input activity patterns of stationary spots	116
Ca ²⁺ dependent model of synaptic plasticity	116
STDP _{sub} model as an implementation of the Butts et al. BTDP rule .	118
Constructing a BTDP rule from the [Ca ²⁺] and STDP _{sub} models . . .	119
Results	120
Interaction of retinal waves with long time scale STDP rules	120
Refinement of receptive fields by retinal waves and STDP with long time scales	124
STDP rules with short decay times that yield burst-timing dependent plasticity	126
The interaction of travelling waves with the STDP _{sub} model	131
Refinement of receptive fields by the STDP _{sub} model	138
Disrupting retinal wave propagation has contrasting effects on RF refinement under two different models of BTDP	141
Discussion	150
6 The influence of delayed inputs on pattern formation and its application the development of simple cells in primary visual cortex	157
Background	157
Methods	161
Spiking neuron simulation with multiple input populations	161
Calculating phase shifts between connectivity patterns	162
Construction of spatiotemporal receptive fields	163
Response of output neurons to test stimuli	164
Direction selectivity index	166
Artificially constructed STRFs	166
Results	167
Outline of synaptic strength dynamics with multiple, delayed input populations	167
Deriving synaptic strength dynamics for two input populations with a delay	172

Development of connectivity patterns from two delayed input populations in a simulation of spiking neurons	176
Synaptic strength dynamics for multiple delayed input populations . .	179
Development of connectivity patterns from four input populations in a simulation of spiking neurons	181
Dependence of direction selectivity on the relative time delays between inputs	184
Discussion	190
7 Conclusions	201
Summary of aims and results	202
Pattern formation with spike-timing dependent plasticity	203
A role for pattern formation in the development of receptive fields	204
Potential roles for retinal waves in retinocollicular refinement	205
Development of selectivity in simple cells	207
8 References	209

Acknowledgements

First and foremost, my most grateful thanks goes to my supervisor, Prof. Wyeth Bair, for his guidance, for teaching me, and for instilling in me a joy for scientific enquiry. Wyeth is unique among teachers, always having to hand simple yet pertinent analogies that paint real pictures of complex and abstract ideas, illuminating the deeper meaning behind the details of an observation. I hope I can carry some of these skills forward in my own scientific endeavours.

I also extend many thanks to my other supervisors, Prof. Andrew Parker and Prof. Paul Bressloff, for helpful and constructive discussions, and for their support and guidance.

Others deserving of more thanks than I can express here, are my college advisor, Kristine Krug, for her advice over the years and for making me feel welcome in her office space; to my good friend, Matthew Broome, who is a constant source of refreshment for the mind, whether up a mountain in the freezing cold, in a closed time-like-curve, or anywhere else for that matter, motivating me to explore our beautiful world; to my friends here at Oxford, in and outside the lab, especially the Bair, Parker and Krug labs for keeping me happy and sane, in particular Pamela for the laughter in multiple continents, Betina for sharing her imagination and Sarah for getting stuck in trees with me.

I am also extremely grateful for financial and educational support from the Wellcome Trust and Oriel College.

Finally, thanks go to my family, for their love, and the strength they have given me to pursue my interests.

Introduction

Travelling waves are a widely reported feature of the developing brain (Feller, 1999; Ben-Ari, 2001; Blankenship & Feller, 2010). The type of travelling wave to be discussed in this thesis is principally characterised by a localised and transient elevation in cellular activity that propagates within a network of neurons. The activity of interest is the generation of action potentials, or spikes, which set into motion a number of downstream cellular events, involving the flux of calcium ions (Ca^{2+}) into neurons, which is considered essential for plastic physiological and structural changes to the synaptic connections within that network (Chklovskii et al., 2004; Markram et al., 2011; Feldman, 2012).

Travelling wave activity patterns have been observed during early developmental stages in a number of brain areas, including the hippocampus (Leinekugel et al., 1998), neocortex (Garaschuk et al., 2000; Conhaim et al., 2010), cerebellum (Watt et al., 2009) and retina (Meister et al., 1991). Remarkably, these phenomena occur spontaneously, largely thanks to the noise-driven and electrically excitable regime in which developing neural circuits exist. Travelling waves, or types of activity that approximate them, may also be induced later on in development, and even throughout life, by sensory experience. For example, when walking through a forest, regions

of bright sky interspersed by dark tree trunks cast high contrast luminance patterns that move across our retinae.

Because travelling waves induce strong correlations in the activity of nearby neurons over short periods of time, and therefore facilitate correlations between those groups of neurons and their postsynaptic targets, they are widely proposed to play an instructive role in the structural development of neural networks via correlation-based mechanisms of plasticity. As Donald Hebb concisely hypothesised:

When an axon of cell A is near enough to excite a cell B and repeatedly or persistently takes part in firing it, some growth process or metabolic change takes place in one or both cells such that A's efficiency, as one of the cells firing B, is increased. - Hebb (1949)

So enduring was this statement, the strengthening of synaptic connections between co-active neurons has come to be known as Hebbian plasticity. However, in this form, Hebb's postulate allowed for the unbounded growth of connection strengths. Such a neural network would become unstable. Several processes that can regulate synaptic strengths, such as homeostasis, synaptic scaling, long-term depression and synaptic redistribution (reviewed in Abbott & Nelson, 2000, and Turrigiano & Nelson, 2004), have since been proposed and worked out. However, the notion of what constitutes coincident activity in Hebb's postulate – whether neurons must be co-active within hundreds of milliseconds, tens of milliseconds, or even the few milliseconds of individual spikes – and a mechanism to detect coincident activity at individual synapses, has been more difficult to formulate, and probably depends on the task at hand in any particular brain area. One leading hypothesis is that individual spikes, and the times at which they are triggered, are the relevant units of computation and plasticity. This idea was theorised by Gerstner et al. (1996), and demonstrated experimentally by Markram et al. (1997). In their experiments,

Markram et al. indicated that coincidence between pre- and postsynaptic activity referred to voltage changes at the postsynaptic terminal of a synapse. These voltage changes are elicited by excitatory postsynaptic potentials (EPSPs), which are evoked by the binding of neurotransmitter to postsynaptic receptors after a presynaptic spike, and postsynaptic action potentials that backpropagate through the dendrites to the synapse. These two features are necessary for the activation of N-methyl-D-aspartate receptors (NMDARs) in the postsynaptic terminal, through which a large Ca^{2+} current can flow into the postsynaptic neuron. Markram et al. (1997) showed that the temporal window in which the requirements for coincidence can be met is on the order of tens of milliseconds. Furthermore, potentiation of the synaptic strength would only take place if the postsynaptic spike followed the presynaptic spike, which captures the causality of Hebb's postulate, whereas depression of the synaptic strength would occur if the temporal order of spikes were reversed. This phenomena came to be called spike-timing dependent plasticity (Song et al., 2000).

Many nuances have since been discovered in addition to the original findings of Markram et al., 1997, as have many other forms of synaptic plasticity that do not depend on coincident activity (Nelson & Turrigiano, 2008; Feldman, 2012). Nevertheless, the principals that underlie the findings of Markram et al. (1997) are conserved at many synapses throughout the brain and form the basis upon which reliable spatiotemporal structures, both in the activities of neural networks and in the surrounding environment, can be represented in the physical wiring of the brain.

The focus of this thesis is therefore to interrogate the interaction between travelling wave activity patterns and spike-timing dependent plasticity (STDP) and to explore the consequences this interaction has for the development of structured connectivity patterns in neural networks. The results presented here have broad application, but

the prevalence of retinal waves during key, early stages of visual system development, and the ubiquity of spatiotemporal structure in our visual environment, which can bias the development of visual responses, has made the developing visual system an ideal arena in which to frame our investigations (Huberman et al., 2008b). In the following, I provide a broad overview of the topics addressed in this thesis, including literature on activity dependent and activity independent development in the visual system, retinal waves, STDP, and computational models of these phenomena. Then, in each results section, I provide a detailed review of the literature that bears more directly on the results presented.

Spontaneous retinal waves

Spontaneous activity in the retina results in the generation of action potentials and Ca^{2+} transients in retinal ganglion cells (RGCs), and is not evoked by light. Indeed, such activity occurs very early in development, before the onset of visually evoked responses in RGCs (Huberman et al., 2008b). With the use of multi electrode arrays and calcium imaging, it became clear that spontaneous retinal activity propagates in travelling waves that sweep across the retina (Meister et al., 1991; Feller et al., 1996). Retinal waves induce strong correlations between nearby RGCs, which become less correlated as the distance separating them increases (Wong et al., 1993; Stafford et al., 2009), thus encouraging the hypothesis that they may act through Hebbian mechanisms of plasticity and instruct, to some degree, the refinement of RGC axons in target nuclei. That retinal waves are a well conserved phenomena adds weight to this proposal, having been recorded from the retinae of mice (Stafford et al., 2009), ferrets (Meister et al., 1991), cats (Meister et al., 1991), rabbits (Syed et al., 2004), macaque monkeys (Warland et al., 2006), chicks (Sernagor et al., 2000) and turtles (Sernagor et al., 2003).

Retinal wave characteristics change over the course of development, the precise

details of which vary between animals, but are generally categorised into three main stages that have been characterised in mice, ferrets and rabbits, and are largely defined by the dominant mechanism underlying wave initiation and propagation (Huberman et al., 2008b). Stage 1 waves occur prenatally in these animals and cover large areas of the retina at slow speeds, requiring gap junction coupling between RGCs (Syed et al., 2004). Stage 2 waves occur during the first one or two postnatal weeks, a period of considerable refinement in retinotopy and eye specificity in downstream targets of RGCs (Torborg & Feller, 2005). Stage 2 waves travel at about half the speed of those in stage 1 (0.1–0.2 mm/s; Feller et al., 1997; Syed et al., 2004; Maccione et al., 2014), and depend on acetylcholine transmission via nicotinic acetylcholine receptors (nAChRs) in starburst amacrine cells (SACs) (Feller et al., 1996; Syed et al., 2004). SACs undergo long refractory periods, lasting tens of seconds to minutes after wave participation, leaving refractory domains that restrict the trajectory along which these waves propagate (Feller et al., 1997; Zheng et al., 2006). Activity in nearby RGCs is strongly correlated by these wave dynamics, and the correlation strength falls off quickly with increasing RGC separation (Wong et al., 1993; Stafford et al., 2009). Furthermore, stage 2 waves elicit activity through all stages of the early visual system, including SC, LGN and V1 (Ackman et al., 2012). Finally, stage 3 waves continue until around the time of eye opening, and coincide with the functional segregation of ON and OFF RGCs (Wong & Oakley, 1996), which respond strongly to either increase or decreases in local luminance, respectively, during visual stimulation. Stage 3 waves depend on glutamatergic transmission from bipolar cells, and cover much smaller distances than in the previous two stages but at similar speeds to stage 1 waves (Syed et al., 2004; Maccione et al., 2014).

Developmental refinement of retinotopy and receptive fields

The early stages of visual development follow a sequence, largely conserved between species, that imparts functional and anatomical specificity to the visual system (Huberman et al., 2008b). In particular, a neuron’s response to visual stimulation becomes progressively more selective for a set of visual features, which characterise its receptive field (RF). Such features might be a circumscribed region of visual space, the orientation of a luminance contour, or the direction of an object’s motion. The theoretical principals developed in this thesis are applied to the development of RFs in the visual system, in particular the refinement of responses to these particular features. We therefore provide a summary of what is understood about the development of such RFs, and its dependence on spiking activity. A concise description of the effects on retinotopic and receptive field development after manipulations to early, spontaneous activity and molecular guidance cues is provided in Table 1.1. In the following, we describe these effects in greater detail.

Establishing retinotopy through molecular guidance cues

The idea that the highly organised arrangement of connections between brain areas is established by molecular cues, which guide growing axons during development, was formally proposed by Sperry in his chemoaffinity hypothesis (Sperry, 1963). The molecular guidance cues responsible for growth in the first stages of the visual system have been well characterised (McLaughlin & O’Leary, 2005), particularly in the superior colliculus (SC) of rodents and chicks. Initially, RGC axons enter the SC through the anterior pole and extend all the way to the posterior pole, far overshooting their correct retinotopic position along the anterior-posterior axis (Simon & O’Leary, 1992a; Simon & O’Leary, 1992b; Yates et al., 2001). Nevertheless, at the correct position along this axis, interstitial branches are sent from the primary axons in a direction perpendicular to it, extending along the medial-lateral axis towards the retinotopically correct termination zone (Yates et al., 2001; McLaughlin et al.,

2003b). Interstitial branching is achieved through the repulsive action of ephrin-As and EphA receptors, each of which are expressed in both the retina and SC in opposing concentration gradients, such that interstitial branching is confined to a small region along the anterior-posterior axis in axons with a particular expression level of ephrin-A and EphA (Yates et al., 2004; Rashid et al., 2005). By this process, a retinotopic map is established between the retina and SC, whereby nearby RGCs in the retina project to nearby areas of the SC thus preserving the representation of visual space. A similar process governs the development of the retinotopic map in the lateral geniculate nucleus (LGN) (Cang & Feldheim, 2013).

Dependence of refinement on spiking activity

Dependence on retinal waves

Retinal waves induce correlations between RGCs that are local both in space and time (Wong et al., 1993). This fact promoted the hypothesis that retinal waves could play an instructive role in the early development of the visual system via Hebbian-like mechanisms of plasticity, which depend on correlated pre- and post-synaptic activity. Such a hypothesis had already been borne out in computational models in which a Hebbian rule was utilised to build retinotopic maps and single neuron RFs in networks of neurons that were driven by input patterns with localised correlations (von der Malsburg, 1973; Bienenstock et al., 1982, Goodhill, 1993). In order to test this experimentally, a number of strategies have been employed to alter retinal wave correlations. One approach has been to artificially induce correlations between RGC axons (Weliky & Katz, 1997; Zhang et al., 2012). Another approach has been to disrupt wave propagation, without abolishing RGC activity altogether (Bansal et al., 2002; Sun et al., 2008a; Stafford et al., 2009). The latter was achieved in many studies by using a genetic mouse model that lacked the $\beta 2$ subunit of nAChRs. Because acetylcholine neurotransmission is an important component of wave propagation during stage 2 waves, retinotopic maps from $\beta 2$ -knockout ($\beta 2^{-/-}$)

mice exhibit considerably perturbed stage 2 waves, which are characterised by faster wave speeds, wide wavefronts and, most importantly, relatively weak correlations between nearby RGCs and relatively strong correlations between distant RGCs, when compared with wild-type counterparts. These manipulations have been shown to affect retinotopy and RF development in different ways, depending on the site of interest along the visual pathway. The major findings for affected development in SC, LGN and V1 are summarised in the following.

Anatomical effects in superior colliculus. Retinocollicular fibres that are sent from a focal point in the retina grow across a large extent of the SC but elaborate and form densely packed arbours in a focal region of SC. In $\beta 2^{-/-}$ mice, however, RGC axons do not elaborate as densely within their appropriate target zone but form arbours over a larger area (McLaughlin et al., 2003a; Chandrasekaran et al., 2005; Pfeifferberger et al., 2006; Dhande et al., 2011). When retinal waves are left intact but instead molecular guidance cues are perturbed, RGC axons do refine dense arbours but in multiple, ectopic locations (Chandrasekaran et al., 2005; Pfeifferberger et al., 2006; Chandrasekaran et al., 2009). These findings suggest that the spatiotemporal structure of retinal waves, and not the mere presence of spiking activity in the retina, play a role in retinocollicular refinement. In contrast to these results, Zhang et al. (2012) reported that artificially inducing correlations across the whole retina in mice had little effect on the organisation of RGC axons in the SC, which elaborated arbours in precise retinotopic locations. To stimulate the retina, Zhang et al. used mutant mice in which RGCs expressed channelrhodopsin, which depolarises the cell under optical illumination. Because channelrhodopsin is not expressed until postnatal day 8 in these mice, it is likely that the majority of retinocollicular refinement had already taken place before optical illumination was effective in stimulating the retina.

Functional effects in superior colliculus. Under normal development, RF diam-

eters in SC more than double in $\beta 2^{-/-}$ mice when compared to wild-type (Chandrasekaran et al., 2005). However, in mice with perturbed guidance cues, RFs were not drastically changed, although some exhibit multiple subfields (Chandrasekaran et al., 2009), reflecting the ectopic anatomical targeting of RGC axons.

Anatomical effects in lateral geniculate nucleus. As was found in SC, the LGN in $\beta 2^{-/-}$ mice receives poorly targeted RGC axons that elaborate over relatively large areas (Grubb et al., 2003). Mice in which waves were intact but guidance cues were disrupted exhibited multiple, ectopic RGC termination zones that were nevertheless refined and densely elaborated (Pfeiffenberger et al., 2006). Multiplicity in the termination zones only appeared along one axis in the LGN, along which the nasal-temporal axis of the retina is mapped.

Functional effects in lateral geniculate nucleus. With the anatomical defects described above, it is surprising that the size of LGN RFs is mostly unaltered in $\beta 2^{-/-}$ mice (Grubb et al., 2003). However, RF centres are distributed over a larger area around their correct retinotopic location (Grubb et al., 2003, Cang. et al., 2008). This loss of retinotopy is more severe for mapping of the nasal-temporal axis of the retina rather than for the ventral-dorsal axis.

Anatomical effects in primary visual cortex. Focal injections of retrograde tracers in cortex label a set of cells in LGN that are approximately 50% more dispersed in $\beta 2^{-/-}$ mice than in wild-type mice (Cang et al., 2005), suggesting that the disruption in retinotopy is compounded from LGN to cortex in $\beta 2^{-/-}$ mice.

Functional effects in primary visual cortex. The defects in anatomical retinotopy were reflected in the functional retinotopy recorded from the cortex of $\beta 2^{-/-}$ mice using intrinsic imaging. The scatter of RFs across the visual field in $\beta 2^{-/-}$ mice was double that of wild-type controls, being most pronounced along the axis that maps nasal-temporal retina (Cang. et al., 2005). This likely reflects, to some de-

gree, the scatter in LGN RFs in $\beta 2^{-/-}$ mice, as individual RFs were no larger in the mutant mice than in controls (Wang et al., 2009). Furthermore, Weliky & Katz (1997) used cuff electrodes to stimulate the optic nerve in ferrets with periodic pulses, starting one week before eye opening and lasting two weeks, thus artificially inducing correlations across all RGC axons. Animals that received this treatment exhibited impaired orientation and direction selective responses in V1 relative to controls. However, this stage of development coincides with the development of distinct ON and OFF activity patterns in the retina (Wong & Oakley, 1996), making it difficult to interpret this effect as a result of disrupted ON and OFF correlations or disrupted retinal wave correlations. Moreover, the particular pattern of activity induced by the pulse stimulation, if any, is not known.

Dependence on visual experience

The visual environment is highly correlated in space and in time (Dong & Atick, 1995). This, and the finding that constraining patterned vision can reliably determine response properties in the visual cortex (Blakemore & Cooper, 1970), strongly motivated the hypothesis that visual experience plays an instructive role in the development of visual circuits, and this can be achieved through a Hebbian mechanism of plasticity (von der Malsburg, 1973). Since the early studies of Hubel and Wiesel (Hubel & Wiesel, 1963), evidence has amassed regarding the ability of visual experience to influence the development of cortical response properties (reviewed in Huberman et al., 2008 and Espinosa & Stryker, 2012). Nevertheless, whether visual experience is instructive or only permissive for cortical development is still largely debated. Likewise, in the LGN and SC, the precise role of visual experience is unclear. Visual deprivation experiments, in which eyelids were sutured, suggest that the LGN develops mature functional responses independent of visual experience (Wiesel & Hubel, 1963; Blakemore & Vital-Durand, 1986b). However, this has more recently been called into question, by the finding that visual experience is possible

through closed eyelids (Krug et al., 2001). When reared in the dark, response properties in the LGN fail to mature (Akerman et al., 2002). In the following, current ideas surrounding the role of visual experience in visual development is provided for the SC and LGN, which is summarised in Table 1.2, and for the cortex, which is summarised in Table 1.3.

Effects in the superior colliculus and lateral geniculate nucleus. Many of the early studies into the role of visual experience in geniculate and collicular development suggested that response properties in these areas were innately determined, or at least that they depend on spontaneous activity prior to normal vision. Under normal rearing conditions, RF diameters in the ferret LGN undergo considerable refinement during the first weeks of visual experience, decreasing in size by 50–80% (Tavazoie & Reid, 2000). A similar developmental trend is seen in cats (Daniels et al., 1978). However, experiments in cats (Hubel & Wiesel, 1963) and in monkeys (Blakemore & Vital-Durand, 1986b), in which the eyelids are sutured closed to prevent early visual experience, suggested that such experience had little affect on the response properties of cells in the LGN. These cells had spatial configurations that matched those recorded in normally reared animals, whereby the centre of LGN RFs responded well to either increases or decreases in luminance (called ON and OFF responses, respectively), and were suppressed by the same luminance changes in the RF surround. In the SC of cats, responses that are selective to the direction of a moving visual stimulus are found for all directions (Cynader et al., 1975). However, rearing kittens from eye opening in a world consisting of motion in just one direction had no effect on the distribution of directions to which SC cells were selective (Cynader et al., 1975).

Whether suturing the eyes is a suitable way of preventing vision was called into question by a study in which visual responses in the LGN and V1 were recorded

through closed eyelids in ferrets, prior to the normal age at which they open (Krug et al., 2001). Moreover, preventing vision in ferrets by dark rearing, instead of suturing, prevents the normal development of selective ON or OFF responses in the LGN (Akerman et al., 2002). Rather, many LGN cells in these animals exhibit both ON and OFF responses.

Ferrets are a particularly useful model, as they are born seven weeks before eye opening, exhibiting visual responses two weeks before eye opening. This late onset of vision after birth makes them particularly useful for studying the development of visual function around eye opening, as neuronal responses are more robust than in the cat and monkey, which are born at a later developmental stage. Rodents, too, share this advantage. A summary of visual development in cats, monkeys, ferrets and rodents, under normal rearing conditions and under visual deprivation, is provided in Table 1.2.

More recent studies in rodents have added to the results of Akerman et al. (2002) in ferrets to show that visual experience can play a role in the development of the LGN and SC, but only for the maintenance of RF properties that first develop independent of vision. In rodent SC, RFs reach their adult-like size a day or so before eye opening (Fortin et al., 1999; Carrasco et al., 2005). This is seen even in dark reared animals, which is consistent with the dependence of the refinement of RFs in SC upon retinal waves at an earlier stage (Chandrasekaran et al., 2005). However, in hamsters that are dark reared beyond 80 days of age, RF diameters are seen to expand by as much as 100%, suggesting that continued visual experience is required for maintenance of mature RFs (Carrasco et al., 2005). In the LGN of mice, LGN cells receive approximately 20 retinal inputs before eye opening. These are pruned over the course of a couple of weeks after eye opening to just 2 or 3 inputs, and the surviving synapses increase in strength considerably over that time (Chen & Regehr,

2000; Hooks & Chen, 2006). Chronic dark rearing appears to have only a modest effect on synaptic remodelling in mice. However, a few days of visual experience before is enough to set off a short period of heightened retinogeniculate plasticity. After this initial visual experience, four or more days of dark rearing significantly increases the number of retinal inputs to LGN cells, and the strength of these inputs decreases, an effect that is reversible with restored vision (Hooks & Chen, 2006; Hooks & Chen, 2008). These results add to the hypothesis that experience acts only to maintain previously refined RFs in the LGN and SC. However, it has not yet been demonstrated whether visual experience is instructive or merely permissive for normal RF development.

Effects in the primary visual cortex. A similar discrepancy regarding the role of early experience in visual development is seen between several studies in V1, which are summarised in Table 1.3. Cats have been a model animal for studying visual cortical development, but their young age of 8 days at eye opening, and the weak, habituating responses of cortical neurons at this age, make it difficult to interpret recordings of single neurons and their selectivity for different features. Hubel & Wiesel (1959) first showed that many neurons in V1 have RFs consisting of alternating, striped regions that were either ON or OFF responsive. Such cells are therefore particularly sensitive to elongated stimuli, such as a bright bar on a dark background, but only if the bar was presented at a particular orientation. In a later study, Hubel & Wiesel showed that such orientation selectivity was well established at eye opening (Hubel & Wiesel, 1963), though not as well developed as in adults. This was corroborated by Braastad & Heggelund (1985). It is possible, however, that these studies suffer from a sampling bias, in that neurons with robust, selective responses were more likely to be discovered than weak, nonselective neurons. A much larger sample in Blakemore & van Sluyters (1975) showed that orientation selectivity is very weak at eye opening in cats. However, there is consensus that

orientation selectivity in V1 increases through early visual experience and that, by 8 weeks of age, it has reached adult-like maturity (Blakemore & van Sluyters, 1975; Braastad & Heggelund, 1985; DeAngelis et al., 1993a).

The development of orientation selectivity in cats can be disrupted by preventing normal visual experience via dark rearing or by suturing closed the eyelids (Blakemore & van Sluyters, 1975; Sherk & Stryker, 1976; Braastad & Heggelund, 1985), though these reports differ in the amount of disruption caused. Nevertheless, it is clear that correlations in the visual environment determine how V1 develops. When kittens are reared in artificial scenes consisting of high contrast stripes that are oriented along the same axis, orientation selectivity in V1 is strongly biased for that same orientation (Blakemore & Mitchell, 1973; Blakemore & van Sluyters; Blakemore & Cooper, 1970; Stryker et al., 1978). However, these results are again blurred by discrepancies between the percentage of neurons found to be influenced by the patterned environment in which the kittens were raised. Blakemore & Cooper reported that nearly all neurons encountered acquired selectivity for the orientation specified by the environment, whereas only half did so in the study of Stryker et al. (1978).

These issues in the cat may have been a result of unreliable responses due to anaesthesia used. This is less of a problem in ferrets (Chapman & Stryker, 1993), which have become a prominent model of visual cortical development. Before the eyes open naturally, orientation selectivity in ferrets is poor, but does begin to improve immediately, reaching adult like levels by postnatal week 7 (Chapman & Stryker, 1993; Li et al., 2006). As in cats, visual deprivation by dark rearing or suturing the eyelids prevents the full maturation of orientation selectivity (Chapman & Stryker, 1993; Li et al., 2006).

Manipulation of retinal activity				
	Artificially induced activation of the whole retina	Disrupting retinal wave propagation by blocking nicotinic ACh transmission in the retina	Transgenic expression of the $\beta 2$ subunit nicotinic ACh receptor in the retina of $\beta 2$ -knockouts	Manipulation of molecular guidance.
Influence on retinal activity patterns	RGC axons across the whole retina were stimulated either by electrodes that surrounded the optic nerve [1], or by optogenetic stimulation [2]. In both cases, stimulation is effective during stage 3 waves. However, the <i>in vivo</i> patterns of retinal activity due to this stimulation are unknown.	Retinal waves were disrupted by pharmacological inhibition of nicotinic ACh receptors, or by genetic knockout of the $\beta 2$ subunit containing nicotinic ACh receptors ($\beta 2$ -knockouts). Both approaches reduce correlations between nearby RGCs, and increase correlations between distant RGCs, relative to controls [3-5]. In $\beta 2$ -knockout mice, stage 2 waves are faster [5].	Stage 2 retinal waves are restored in these mutant mice. The waves are spatially restricted, but are otherwise very similar to wild-type waves [6].	Stage 2 waves are slightly broader in ephrin-A2, ephrin-A3 and ephrin-A5 triple knockout mice, and induce strong correlations between RGCs over greater separations than in controls [7].
Anatomical changes in superior colliculus	Optogenetic stimulation has little effect on retinotopic organisation of RGC afferents [2]. This may be because stimulation was only possible after the main stage of retinocollicular refinement by stage 2 retinal waves.	RGC termination zones are significantly larger than in controls, severely disrupting retinotopy [6, 8-10]. In adult mice, retinotopic organisation is no more precise than in P4 controls [8]. However, a coarse retinotopy is still preserved [6, 8-10]. Retinotopy along the nasal-temporal axis of the retina is disrupted the most [9]. Single RGC axon arbours are significantly larger than in controls [11].	RGC termination zones cover significantly larger areas than in controls, but only in binocular regions of SC. Termination zones in monocular regions are as refined as in controls [6].	Knockout of ephrin-A2, ephrin-A3 and ephrin-A5 in mice results in multiple, ectopic termination zones from nearby RGCs [10]. The same effect is observed in mutant mice in which BMP is over expressed in the retina [9].
Functional changes in superior colliculus	Unknown.	RF diameters in adult mice are on average twice as large as those in controls. RFs are particularly elongated along the azimuth, in agreement with anatomical defects [9].	Unknown.	On average, RFs are only slightly larger in adult BMP mutant mice, though several RFs are substantially larger, and even contain multiple subfields [12]. Subfield separation is consistent with the distance separating the multiple RGC termination zones in these mutants.
Anatomical changes in lateral geniculate nucleus	Unknown.	RGC termination zones are larger than in controls, but coarse retinotopy is maintained [6, 10, 13]. However, arbours of single RGC axons are not significantly larger than in controls [11].	RGC termination zones are disrupted in the same way as in SC, in a manner that also depends on whether the termination zones lie in binocular or monocular regions of the LGN [6].	In ephrin-A2/A3/A5 knockout mice, RGC termination zones occupy larger areas of LGN than in controls, and in a few cases establish multiple, ectopic termination zones [10].
Functional changes in lateral geniculate nucleus	Optic nerve stimulation does not disrupt the functional segregation of ON and OFF responses [1].	The spatial characteristics of individual RFs are not significantly different from those in controls [13]. However, the scatter of RF centres across retinotopic space is increased along the nasal-temporal axis of the retina when compared to controls [13].	Unknown.	Unknown.
Anatomical changes in V1	Optic nerve stimulation has no effect on the spatial organisation of orientation columns [1].	Afferents to localised areas of V1 arrive from larger regions of LGN than in controls [14].	Unknown.	Afferents to local areas of V1 arrive from broader regions of LGN than in controls, but with less retinotopic disruption than in animals that have perturbed retinal waves [15].
Functional changes in V1	Optic nerve stimulation results in a significant reduction in cortical orientation selectivity [1].	Intrinsic imaging and multi-unit recordings show greater retinotopic scatter of RF centres along the azimuthal axis when compared to controls, consistent with findings in the LGN [14]. However, the size of single unit RFs is not significantly different from controls [16].	Unknown.	Intrinsic imaging shows that retinotopy is disrupted in mutant mice that do not express ephrin-A2 or ephrin-A5, but not to the same extent as in mice with perturbed retinal waves [15].

Table 1.1: (*Table on previous page.*) **Summary of manipulations to spontaneous retinal activity and molecular guidance, and the effect such manipulations have on anatomical and functional refinement in the visual system.** [1] Weliky & Katz, 1997; [2] Zhang et al., 2012; [3] Sun et al. 2008a; [4] Sun et al, 2008b; [5] Stafford et al., 2009; [6] Xu et al., 2011; [7] Pfeiffenberger et al., 2005; [8] McLaughlin et al., 2003a; [9] Chandrasekaran et al., 2005; [10] Pfeiffenberger et al., 2006; [11] Dhande et al., 2011; [12] Chandrasekaran et al., 2009; [13] Grubb et al., 2003; [14] Cang et al., 2005; [15] Cang et al., 2008a; [16] Wang et al., 2009. *Abbreviations:* RGC - retinal ganglion cell; ACh - acetylcholine; RF - receptive field; LGN - lateral geniculate nucleus; SC - superior colliculus; BMP - bone morphogenetic protein.

	Visual experience	
	Normal	Dark reared / eyelid suture / restricted visual motion
Cats and Monkeys	The developmental trajectory of LGN RFs in cats comprises a gradual reduction in RF diameter from 1°-30° at P6, just before eye opening, to 1°-5° after 4 weeks of visual experience [1]. This coincides with a gradual increase in surround suppression, which is practically absent at P6 but is very prominent by P34, reducing the RF centre to 0.2°-0.5° in size [1]. The percentage of cells in LGN with selective ON and OFF responses also increases, from 60% at P6 to 80% by P34 [1].	Rearing cats [6] or monkeys [7], after suturing closed their eyelids from birth, has very little effect on spatial RF properties in the LGN, on the whole, bar a few exceptional cells that were larger than normal and lacked surround suppression when tested after 3 months of deprivation [6]. Direction selectivity in the SC is no different in normally reared cats than if they had been reared in a visual environment consisting of motion in just one direction [8].
Ferrets and Rodents	RFs in the LGN of ferrets decrease in diameter from 3°-45° at eye opening to 1°-6° in adults [2]. This is consistent with a reduction in retinal inputs from ~20 at P12 to just 1-3 by P28 in mice [3]. In the rodent SC, RF sizes are already mature in the rat, being ~10° in diameter, and do not change much throughout visual experience [4]. In hamsters, RF diameters undergo refinement from ~35° before P28 to ~20° by P51 [5].	Ferret LGN cells respond to visual stimulation through naturally closed eyelids, prior to eye opening [9]. Dark rearing ferrets prior to eye opening affect the functional segregation of ON and OFF response types in the LGN [10]. Dark rearing does not affect the refinement of the total retinal input to the mouse LGN [11]. However, if the mice receive at least a few days of visual experience before P20, followed by dark rearing for at least 4 days, the number of retinal inputs to individual LGN cells increases [11, 12]. A return to normal levels of innervation is achieved if vision is restored [12]. In the hamster SC, dark rearing has no effect on the refinement of RF size, unless it continues beyond P80, after which RF diameters increase to immature sizes [5]. Thus, visual experience has role in maintenance, rather than refinement, in rodent LGN and SC.

Table 1.2: **The development of the lateral geniculate nucleus and superior colliculus and its dependence on different types of early visual experience.** [1] Daniels et al., 1978; [2] Tavazoie & Reid, 2000; [3] Chen & Regehr, 2000; [4] Fortin et al., 1999; [5] Carrasco et al., 2005; [6] Hubel & Wiesel, 1963; Blakemore & Vital-Durand, 1986b; [8] Cynader et al., 1975; [9] Krug et al., 2001; [10] Akerman et al., 2002; [11] Hooks & Chen, 2006; [12] Hooks & Chen, 2008. *Abbreviations:* RF - receptive field; LGN - lateral geniculate nucleus; SC - superior colliculus.

Visual experience			
	Normal	Dark reared / eyelid suture	Patterned environment
Cat	Some debate remains surrounding the extent of orientation tuning at eye opening. [1] and [2] report that OS is well established just after eye opening, though not as well tuned as in adults. [3] and [4] report that OS is very poor just after eye opening, and improves substantially with visual experience. Nevertheless, there is a general consensus that by week 4, OS in V1 is adult-like, though the sharpness of simple cell tuning continues to improve until about week 8 [2, 5].	OS fails to reach the same sharpness of tuning as in adults when the kitten is reared in the dark [2, 4] or when the eyelids are sutured [4, 9]. As with normal reared kittens, the extent to which OS suffers, because of rearing in the dark or with sutured eye, differs between reports.	For kittens reared in environments consisting of high contrast stripes that are oriented along only one axis, OS in V1 cells is biased towards the same orientation [3, 4, 10, 11]. OS can be rapidly induced in as little as one hour of exposure to the stipe patterns [3]. Reports differ in how many cells acquire the experienced orientation, ranging from the vast majority [10] to only about half of all cells [11]. Exposure to visual environments with motion in one direction encourages a DS bias for the same direction [12, 13]. DS may be reversed if the experienced motion direction is reversed before 4 weeks of age [13]. Exposure to spots encourages LGN like RFs and fewer orientation selective cells to develop. Large spots, which contain locally distinct oriented contours, induce more OS than do small spots [4].
Ferret	Through postnatal weeks 4 and 5, which are before eye opening, OS in V1 is poorly tuned [6, 7] and DS is absent [7]. OS does increase during the weeks before eye opening, possibly due to vision through the eyelids [8], but DS does not develop until after eye opening [7].	Visual deprivation, either by suturing the eyes until 8 weeks of age [6] or by dark rearing [7] prevents the maturation of OS beyond that found in normal ferrets by the time of eye opening. Dark rearing prevents any DS from developing [7].	Exposure to visual motion is necessary for the development of DS, and exposure to motion in a single direction enhances selectivity for that same direction [14, 15].

Table 1.3: **The development of primary visual cortex and its dependence on different types of early visual experience.** [1] Hubel & Wiesel, 1963; [2] Braastad & Heggelund, 1985; [3] Blakemore & Mitchell, 1973; [4] Blakemore & van Sluyters, 1975; [5] DeAngelis et al., 1993a; [6] Chapman & Stryker, 1993; [7] Li et al., 2006; [8] Krug et al., 2001; [9] Sherk & Stryker, 1976; [10] Blakemore & Cooper, 1970; [11] Stryker et al., 1978; [12] Cynader et al., 1975; [13] Daw & Wyatt, 1976; [14] Li et al., 2008; [15] van Hooser et al., 2012. *Abbreviations:* OS - orientation selectivity; DS - direction selectivity; RF - receptive field.

Most striking in the development of V1 is the selectivity for motion in a given direction, which is not present in ferrets until after eye opening and fails to develop when they are visually deprived (Li et al., 2006; Li et al., 2008). As with orientation selectivity, a bias for motion in a particular direction can be induced in cats (Cynader et al., 1975; Daw & Wyatt, 1976) and ferrets (Li et al., 2008; van Hooser et al., 2012), if the animal is only exposed to motion in that direction.

Despite these convincing reports that visual experience has some influence over the development of selectivity in V1, it remains unclear whether the spatiotemporal correlations in the visual environment instructs this development, or whether such environments simply permit the maturation of neurons that have innate, albeit weak and in some cases undetectable, preferences for those spatiotemporal correlations. This will become a major focus of this thesis, and our results suggest tests that may help to resolve this issue.

Experimental findings and models for spike-timing dependent plasticity

The spike-timing dependence of synaptic plasticity, whereby the relative timing of pre- and postsynaptic spikes determine the magnitude and sign of changes in synaptic strength, was predicted by Gerstner et al. (1996) and first demonstrated by Markram et al. (1997). They established the phenomenon as a form of Hebbian plasticity, whereby the local membrane potential at the synapse, being influenced by both EPSPs and backpropagating action potentials, acts as the required associative signal between pre- and postsynaptic activity. The timing dependence of STDP was characterised in even greater detail by Zhang et al. (1998) and Bi & Poo (1998). They showed that changes in synaptic strength decayed exponentially if the latency between pre- and postsynaptic spikes increased over just tens of milliseconds, signifying that the membrane voltage at the synapse might underlie plasticity.

To understand how STDP may impact changes in connectivity within a network of neurons, Kempter et al. (1999) modelled a STDP rule, based on the exponential decay observed by Zhang et al. (1998). The Kempter et al. model registered all pairs of pre- and postsynaptic spikes, which contributed linearly to the overall change in synaptic strength. Under the assumption that these changes occurred slowly, Kempter et al. derived an important mathematical description of synaptic plasticity, with which changes in the connectivity of the network could be approximately calculated based on known patterns of spiking activity within the network. Equipped with this idea, new models were devised to show how STDP could interpret local correlations in the sensory environment to instruct the development of cortical circuits (Song et al., 2000; Song & Abbott, 2001).

The theoretical results of Kempter et al. (1999), powerful in their simplicity, were not able to account for a number of complexities borne out in future STDP experiments. STDP could not be described simply as a linear summation of synaptic updates. Rather, the change in synaptic strength due to one pair of pre- and postsynaptic spikes had a non-linear dependence on the recent history of spikes (Froemke & Dan, 2002; Froemke et al., 2006). Today, a plethora of additional variables are now known to modulate STDP, including membrane voltage, intracellular calcium concentration, synapse location, conductance state and neuromodulation (reviewed extensively in Feldman, 2012). Attempts to reconcile these complexities with mathematical formulations of STDP have recently produced more mechanistic models of STDP. Where some account for the voltage dependence of STDP (Clopath et al., 2010), others model the intracellular calcium concentration as the determining variable (Shouval et al., 2002; Cai et al., 2007; Graupner & Brunel, 2012). In particular, Graupner & Brunel (2012) demonstrated that even a relatively simplified description of the Ca^{2+} dependence might underlie the many forms of STDP that have been observed for single spike pairs at different types of synapse.

The Ca^{2+} and dependent models of STDP are of particular relevance for development during retinal waves, which evoke large Ca^{2+} transients throughout the early visual system (Ackman et al., 2012). Moreover, the slow dynamics of retinal waves, which elicit bursts several seconds long, obscure correlations between pre- and post-synaptic activity over the fast time scales usually associated with STDP (Caporale & Dan, 2008; Kempter et al., 1999; Song & Abbott, 2001). However, Butts & Rokhsar (2001) proposed that plasticity rules sensitive to correlations over seconds, rather than tens of milliseconds, would enable retinal waves to instruct refinement in downstream visual areas in a Hebbian manner. A rule matching this description, and is sensitive to the timing of spike bursts, rather than individual spikes, was then found at the developing retinogeniculate (Butts et al., 2007) and retinocollicular synapses (Shah & Crair, 2008). This burst-timing dependent plasticity (BTDP) rule strengthens synapses when pre- and postsynaptic bursts are offset in time by up to ± 0.5 s, and weakens synapses for temporal offsets greater than this.

With the considerable complexity surrounding STDP, there are no models that can reliably account for changes in synaptic strength that result from arbitrary patterns of pre- and postsynaptic activity. The simplicity of the theoretical results in Kempter et al. (1999) therefore make it an appealing framework in which to begin an examination of new STDP-related phenomena. We derive many of our own results within this framework, but also examine other models of timing dependent plasticity, including those presented by Butts et al. (2007) and Graupner & Brunel (2012), and their interaction with travelling waves.

Models of Hebbian plasticity for visual development

Some of the first models of plasticity to capture the essence of Hebb’s postulate provided valuable insights into the role of correlated activity in visual development.

In particular, strong correlations between nearby presynaptic neurons were shown to play an instructive role in the development of RF structures and retinotopic maps seen throughout the visual system (Willshaw & von der Malsburg, 1976; von der Malsburg, 1973; Linsker, 1986a; Linsker, 1986b; Goodhill, 1993). These formulations of Hebbian plasticity required an additional mechanism to impose competition between synapses, without which synapses would undergo unbounded strengthening (Miller & Mackay, 1994; Miller, 1996), but it is difficult to pin down a biological analog for the mechanisms that were proposed. Experimentally observed STDP, however, provides a built in mechanism for synaptic competition, as post- followed by presynaptic spikes leads to synaptic weakening. At a number of different synapses, weakening can also occur for other spike sequences (Caporale & Dan, 2008).

Song & Abbott (2001) utilised the intrinsically competitive nature of STDP to model how RFs and topographic maps could emerge in networks when spatially and temporally correlated activity was present in the inputs. Likewise, the competitive nature of BTDP has been shown to drive refinement. During stage 3 retinal waves, ON and OFF RGCs participate in waves with a 1 s delay (Kerschensteiner & Wong, 2008). Gjorgjieva et al. (2009) demonstrated how these temporal correlations could engage BTDP to drive the segregation of ON and OFF inputs to LGN.

The temporal asymmetry common to STDP rules observed in the neocortex (Feldman, 2000) provides additional sensitivity to temporal correlations and causality in neural networks, which has been exploited in models to address the development of direction selective responses in mammalian cortex (Hubel & Wiesel, 1962; Rao & Sejnowski, 2003; Wensch et al., 2005). In these models, networks were trained with visual stimuli moving in a biased direction, driving wave-like patterns of activity in a layer of sensory input neurons. By engaging the temporal asymmetry of STDP, these motion stimuli were able to build an asymmetry into recurrent connections

in a layer of output neurons that were driven by the inputs. Thus, after training, a neuron located at the centre of the output layer received strong excitation when stimuli moved in the same direction as the training stimulus, but weak excitation when stimuli moved in the opposite direction, endowing it with direction selectivity. A similar approach was also shown to yield direction selective neurons by building asymmetries into feedforward networks (Buchs & Senn, 2002; Honda et al., 2011). However, whether these connectivity motifs are present in the cortex is yet to be determined. Experimental data points towards a different model of direction selectivity (Adelson & Bergen, 1985), whereby spatially offset inputs with different response latencies reflect a well known motion-detection circuit (Hubel & Wiesel, 1962; Saul & Humphrey, 1990; de Valois & Cottaris, 1998; De Valois et al., 2000; Saul et al., 2005). It has yet to be shown whether STDP can be utilised to build such a circuit.

Open questions and aims of this thesis

Travelling wave phenomena in the brain have incited much interest due to the correlations they induce, which are localised both in space and in time, thus promising a potential role for them in instructing neural circuit development through Hebbian mechanisms of plasticity. Indeed, this has motivated many models and experiments, as reviewed above, that demonstrate the effectiveness of localised correlations in performing this task and instructing the refinement of neural circuits, in particular in the visual system. However, formal analyses of these models tend to concentrate on the importance of spatial correlations alone (Kempster et al., 1999; Clopath et al., 2010; Gjorgjieva et al., 2011), or temporal correlations alone (Gjorgjieva et al., 2009). In other words, correlations have been treated as space-time separable, and analysed thereon. This thesis is focused on understanding how space-time inseparable correlations, like those associated with travelling waves and stimulus motion, could engage STDP rules to build structure into neural networks.

In Chapter 3, we invoke a simple but plausible network model to examine the interaction between travelling waves and STDP. We derive mathematical expressions that describe this interaction in one-dimensional networks and the influence it has on the synaptic connectivity in the network, and find analytical solutions for these equations. These analytical results reveal a novel function for STDP, whereby its interaction with travelling wave activity can drive pattern formation in the synaptic strengths and, as such, build periodic structures into the connectivity of neural circuits, which we verify using simulations of spiking neurons. This type of pattern formation falls into a large class of pattern forming processes that are found throughout biology (Turing, 1952; Meinhardt & Gierer, 2000), and provide further theoretical grounding to results in previous STDP models (Buchs & Senn, 2002; Honda et al., 2011). We demonstrate critical roles for the wave speed and time scales of STDP in determining the structure of the connectivity pattern. Using these results, we assess the applicability of pattern formation to travelling wave phenomena in several areas of the brain, and stress its relevance to development in the visual system. We use further simulations to examine the robustness of pattern formation to a variety of noise sources that are common in developing neural circuits, and argue that pattern formation holds up well with levels of noise observed experimentally.

In Chapter 4, we use simulations of spiking neurons and numerical solutions to the equations derived in Chapter 3 to demonstrate how pattern formation by STDP and travelling waves can influence the development of receptive fields (RFs). We demonstrate how the wave speed, STDP time scale and initial RF size, together determine the structure of RFs that develop when driven by travelling waves. Refinement of RFs by the pattern formation process is markedly different from refinement in previous models of Hebbian plasticity (Willshaw & von der Malsburg, 1976; Goodhill, 1993; Miller & MacKay, 1994; Song & Abbott, 2001) in that it does not require

correlations in the inputs to decay over space. We further demonstrate how two-dimensional RFs can be refined by the same pattern formation process, even when waves follow complex and noisy trajectories across the input layer. These results are particularly relevant to the development of RFs in the visual system, and yield new predictions that can be tested using current experimental paradigms.

In Chapter 5, we examine a Hebbian role for retinal waves in the development of RFs in superior colliculus. We use simulations of one-dimensional networks of spiking neurons to compare two implementations of STDP that are appropriate for the slow dynamics of retinal waves, and that are also constrained to yield the experimentally observed BTDP rule at retinofugal synapses (Butts et al., 2007). The major difference between these two implementations of STDP is the manner in which synaptic weakening depends on correlated activity. This difference has a critical influence on the ability of the each STDP implementation to produce the periodic connectivity patterns that are central to this Thesis. Nevertheless, we show that both STDP rules are able to refine RFs, but that travelling waves are not a necessary requirement for one of the STDP rules. Using simulations of two-dimensional retinal waves, we demonstrate how these two STDP rules have very different consequences for RF refinement when the propagation of retinal waves is disrupted. We discuss how these results relate to recent experimental findings in which retinal wave propagation is stunted yet retinocollicular refinement is unperturbed (Xu et al., 2011). Our results yield predictions regarding the type of synaptic plasticity that operates at stages of development when retinal waves are prevalent.

In Chapter 6, we extend our theoretical work in Chapter 3 to examine the interaction between STDP and travelling waves in networks comprising multiple, delayed input populations, each of which establishes its own connectivity pattern with a shared output neuron. This network architecture resembles the geniculocortical circuit, in

which neurons in the LGN can be categorised by the latency of their responses to visual stimuli (Saul & Humphrey, 1990), with the goal in mind of applying our results to the development of direction selective simple cell RFs in cortical area V1. We first show that the connectivity pattern for each input population acquires the periodic structure demonstrated in Chapter 3. Furthermore, we show that the difference in response latencies between any two input populations determines the relative configuration in space of their respective connectivity patterns. We demonstrate how such spatial offsets between the connectivity patterns, and temporal offsets between their respective input responses, provide the two essential ingredients to endow the output neuron with selectivity for motion in a particular direction, as proposed by Adelson & Bergen (1985). We compare our model with previous models of direction selectivity development (Wimbauer et al., 1997a; Buchs & Senn, 2002; Wenisch et al., 2005; Honda et al., 2011) and point out several possible advantages to our framework. Our results are particularly applicable to experiments in which direction selectivity in V1 can be trained by exposing an animal to moving visual stimuli (Cynader et al., 1975; Li et al., 2008; van Hooser et al., 2012), and they make new predictions about the properties of simple cell RFs that would develop under particular manipulations to the visual stimulus, which may help to determine the role of visual experience, whether instructive or permissive, in cortical development.

General Methods

Overview

We model a 2-layer feedforward network consisting of a presynaptic layer of model input neurons that connect via feedforward excitatory synapses to a single output neuron. In general, the input layer consists of either a 1 dimensional (1D) chain or a 2D sheet of model input neurons, but the actual size of the input layer depends on the particular question being addressed, and therefore is different in different sections of this Thesis. These more specific details are provided in the relevant Methods sections in each chapter. Here, we present the fundamental components of our model that are present throughout the Thesis. An important feature of the model that is common to all chapters is the pattern of activity in the input layer: travelling waves of activity recruit the input neurons, which discharge bursts of spikes and provide excitatory input to the output neurons. How these input patterns are produced, and the way in which they drive activity in the output neuron, is described below.

Spiking neuron simulation

Output activity and synaptic plasticity

The output neuron is simulated using a time step of 1 ms. During each step, a number of processes are simulated that ultimately determine whether the output will generate a spike and, if it does, how this affects the strength of its synaptic inputs. First, input neurons generate spikes from one of the two wave models described below. Each input spike elicits an excitatory postsynaptic potential (EPSP) in the output neuron, which transiently increases an activation variable, $\lambda(t)$, which scales the probability with which the output neuron spikes and thus its firing rate. The EPSP in our model therefore corresponds to the membrane voltage of real neurons.

The EPSP transient that is generated by the n^{th} input spike is described by a difference of exponentials

$$\epsilon_n(t) = \frac{1}{\tau_d - \tau_r} \left[\exp\left(-\frac{(t - t_n)}{\tau_d}\right) - \exp\left(-\frac{(t - t_n)}{\tau_r}\right) \right], \quad (2.1)$$

where $\tau_d = 5 \text{ ms}$ and $\tau_r = 1 \text{ ms}$ are the EPSP decay and rise times, respectively, t is the current time, t_n the input spike time and $\epsilon_n(t) = 0$ for $t - t_n < 0$. The firing rate, $\lambda(t)$, of the output neuron is proportional to its summed EPSPs such that

$$\lambda(t) = r \sum_j \sum_n w_j(t) \epsilon_n(t), \quad (2.2)$$

where w_j is the synaptic strength from input neuron j to the output neuron and is a dimensionless variable. The coefficient r is a constant, chosen to constrain the output neuron to realistic firing rates in the range 10–100 Hz, and has units of spikes per second per volt, in keeping with the analogy $\epsilon(t)$ has with the membrane voltage. At each time step, output spikes are generated by a Bernoulli process. That is to say, a spike is generated if a number, chosen randomly between 0 and 1, is less than or equal to $\lambda(t)$.

After the spike generating process, synaptic strengths are updated using a spike-timing dependent plasticity (STDP) rule. Two rules are used, one asymmetric (Zhang et al., 1998), K_a , and the other symmetric in time, K_s (Fig. 3.1C):

$$K_a(\Delta t) = \begin{cases} \eta A_+ \exp(\Delta t / \tau_+) & \text{if } \Delta t < 0 \\ -\eta A_- \exp(-\Delta t / \tau_-) & \text{if } \Delta t > 0 \\ 0 & \text{if } \Delta t = 0, \end{cases} \quad (2.3)$$

$$K_s(\Delta t) = \eta \left[A_+ \exp\left(-\frac{1}{2} \left(\frac{\Delta t}{\tau_+}\right)^2\right) - A_- \exp\left(-\frac{1}{2} \left(\frac{\Delta t}{\tau_-}\right)^2\right) \right]. \quad (2.4)$$

Here, $\Delta t = t_{\text{in}} - t_{\text{out}}$ is the time difference between spikes belonging to input and output neurons, respectively. The variables τ_+ and τ_- scale the temporal window size for synaptic strengthening and weakening respectively, A_+ and A_- scale the relative degree of strengthening and weakening, respectively, and η is a constant that scales the rate of change in synaptic strengths. Every time an input (output) neuron spikes, all output (input) spikes in the previous $5\tau_-$ seconds are used to update a synapse. When these processes are complete, the simulation is moved to the next time step. Unless otherwise stated, all asymmetric STDP rules are constructed such that $\tau_- = 2\tau_+$, and all symmetric STDP rules are constructed such that $\tau_- = 1.6\tau_+$. This reflects the widely reported finding that synaptic weakening occurs over longer time scales than does strengthening (Caporale & Dan, 2008). Throughout this Thesis, the time scale of the STDP rules is therefore almost always referred to by a single variable, τ_+ , as τ_- simply scales linearly with it. In all simulations, asymmetric rules are constructed using $A_+ = 1.0$ and all symmetric rules are constructed with $A_+ = 3.2$. However, the A_- is set to different values in different sections to adjust the relative degree of synaptic weakening to strengthening. In Chapter 3, all simulations use $A_- = 0.51$ for the asymmetric STDP rule and $A_- = 2.1$ for the symmetric STDP rule. In Chapter 4, the asymmetric STDP rule is constructed using $A_- = 0.55$ for the application to RF refinement, and the symmetric STDP rule uses $A_- = 2.1$. In Chapter 5, $A_- = 0.51$ for the asymmetric STDP rule. The symmetric STDP rule is used for the majority of simulations in Chapter 5: to compare simulations with theoretical results in 1D networks, $A_- = 2.1$; to address receptive field refinement in 1D networks, $A_- = 2.1$; to address receptive field refinement in 2D networks, $A_- = 2.2$. In Chapter 6, $A_- = 0.51$ for the asymmetric STDP rule and $A_- = 2.1$ for the symmetric STDP rule.

Travelling waves in the input layer

1D and 2D plane waves

The input layer comprises point neurons, arranged $20\text{ }\mu\text{m}$ apart, corresponding to the close packing of cell somas, in a 1D chain or 2D square lattice. Plane waves were made to travel the entire length of the input layer. Waves were interleaved with a silent period lasting at least 5s to allow the postsynaptic activation to decay to zero before the next wave traversed the input. The only simulations in which waves were not interleaved with a silent period were those in which the presence of multiple waves was tested (Chapter 3). When the wavefront reached an input neuron, the input neuron underwent a burst of spikes. In time steps of 1 ms, spikes were generated using a Bernoulli process at a given firing rate. In some simulations, input neurons were given background firing rates. These are stated in the relevant results sections. Otherwise, when an input neuron was not being recruited by a wave, it was silent. A schematic of the 1D network is provided in Fig. 2.1A. In all cases, we repeated trials of plane wave simulations with the same model parameters 16 times, using different seeds for the random number generators, such that each trial would produce different spike trains. In simulations where a parameter that was normally fixed was instead a random variable, such as the wave speed, or wave direction, the random number generator for that random variable was reseeded for repeated trials. If none of the parameters were treated as random variables, the random number generator for spike generating Bernoulli process was reseeded. In all cases, original spike sequences were always generated for each trial.

Complex retinal wave patterns

To generate more realistic retinal waves in 2D, we implemented a simulation based on a previously published model of the developing mammalian retina (Feller et al., 1997). The model comprises a layer of starburst amacrine cells (SACs) and a layer of retinal ganglion cells (RGCs). Both layers are arranged on 64×64 square lat-

tices with a cell spacing of $34\text{ }\mu\text{m}$. SACs receive excitatory synaptic input of equal strength from other SACs that are within a radial distance of $120\text{ }\mu\text{m}$. RGCs also receive excitatory input from SACs over the same distance but do not provide synaptic input to other RGCs or SACs. None of the synapses within the model retina undergo synaptic plasticity.

Every cell in the retina model is characterised by a dynamic excitation variable, X_i , which is analogous to its membrane potential, for $i = 1\dots N$ where N is the number of SACs or RGCs. The simulation evolves in time steps of 100 ms and X_i is updated for cells in the SAC (S) and RGC (R) layers according to

$$\frac{d}{dt}X_i^{\text{S}} = n_i^{\text{S}} - \frac{X_i^{\text{S}}}{\tau^{\text{S}}}, \quad (2.5)$$

$$\frac{d}{dt}X_i^{\text{R}} = n_i^{\text{S}} - \frac{X_i^{\text{R}}}{\tau^{\text{R}}}, \quad (2.6)$$

where n_i^{S} is the number of SACs that are active at time, t , and connected to cell i . The maximum possible value of n_i^{S} is therefore 37, as determined by the maximum distance over which two SACs are connected. The activation decay time constants, τ^{S} and τ^{R} , are both 0.1 s . A SAC becomes active when X_i^{S} exceeds a threshold, $\theta_{\text{S}} = 6.0$. Synaptic transmission persists for 1 s , after which the SAC becomes refractory and X_i^{S} is held at zero. Every time a SAC becomes refractory, the refractory period is chosen independently from a normal distribution with mean τ_{ref} (which in Chapter 5 was varied between conditions) and standard deviation $\sigma_{\text{ref}} = 20\text{ s}$. After the refractory period, X_i^{S} can once more be influenced by synaptic input from other SACs. Electrophysiological recordings of stage II waves show that SACs are highly excitable and spontaneously initiate retinal waves (Zheng et al., 2006). We model spontaneous excitation in subthreshold SACs at each time step by increasing X_i^{S} to θ_{S} with a probability, $P_{\text{n}} = 0.0035$, to provide realistic wave initiation rates.

When X_i^{R} exceeds a threshold, $\theta_{\text{R}} = 10.0$, a RGC generates spikes at a rate of

10 Hz for as long as $X_t^R > \theta_R$. The spikes are generated using a Bernoulli process with a temporal resolution of 1 ms. The RGCs in this model retina then act as the input neurons in our 2-layer simulation.

In two sets of simulations, we artificially varied the wave speed without changing any parameters in the retinal wave model but, rather, manipulating the spike times it produced. To do this, we isolated every burst of spikes for each neuron and shifted them in time by scaling the burst onset time by a factor, F_t . In this way, the shape of the wave trajectory, the total number of spikes and the inter-spike intervals within bursts were all preserved. Although the inter-burst interval was decreased, the vast majority remained longer than the temporal width of the widest STDP rule. The retinal wave model is only used in Chapters 4 and 5, and the relevant values for F_t are provided in the Methods sections of those chapters.

We repeated trials of retinal wave simulations with the same model parameters three times, using different seeds for the τ_{ref} random number generator, such that each trial would produce a different sequence of wave patterns.

Computing numerical solutions

Concordance of the spiking model with analytical results is verified by solving Eq. 3.6 in Chapter 3 numerically. Synaptic strengths are initialised in the same way as in the spiking simulation. However, as there is no spiking noise when solving Eq. 3.6 numerically, low amplitude Gaussian white noise is added to the initial synaptic strengths. Nonlinearities are also included by restricting synaptic strengths to the range $[0, 1]$ and by reflecting $\kappa(x)$ (Eq. 3.6) in the y-axis after each convolution to match the alternating direction of waves in the simulation. η was varied depending on the values of τ_+ and v so that $\partial_T w$ did not become chaotic. Finally, when

numerically solving Eq. 3.6 with initial conditions that support a receptive field, an arbour function was applied to prevent synaptic strength dynamics outside of the arbour. The arbour function was implemented by multiplying $A(x)$ with $\partial_T w(x, T)$, where

$$\begin{aligned} A(x) &= 1, \text{ if } |x - x_0| \leq a_0 \\ &= 0, \text{ if } |x - x_0| > a_0 \end{aligned}$$

Analysis of complex wave patterns

Simulated RGC spiking activity was converted into firing rate movies, $M(x, y, t)$, for analysis. Each movie has dimensions $N_x \times N_y \times N_t$, where the $(i, j, k)^{\text{th}}$ bin in M contains the number of spikes fired by a single RGC, at location (x_i, y_j) , during a 100 ms period, t_k . We analysed waves from a section of each simulation lasting 5000 simulated seconds. This constituted analysis of over 4 hours of simulated retinal wave activity from the three trials used for each set of parameters.

During the retinal wave simulation, multiple waves were occasionally present at the same time. Although concurrent waves might collide into or split from each other, for the most part of their duration, they were spatially or temporally distinct. Thus, we sought to isolate waves to accurately measure their individual properties, such as speed and size. Waves were isolated using a centre of mass tracking algorithm. First, the firing rate movie $M(x, y, t)$ was smoothed using a Gaussian filter with a spatial standard deviation (SD) of $34\mu\text{m}$ and temporal SD of 100 ms. After smoothing, pixel values below 0.1 were set to zero. Units near the edges of the model retina received fewer lateral inputs and were less active than those in the centre. Within the central 44×44 RGCs of the model retina, the total level of spiking activity was comparably uniform. Accordingly, waves with a time-averaged centre of mass

(COM) that resided in the outer 10 neurons were discarded from the analysis. Occasionally, small segments of waves were isolated from larger waves. These were discarded from the analysis by removing waves that covered fewer than 500 space-time bins.

To track the COM of different waves, boundaries were drawn around domains of activity for each time frame in M and the COM calculated for each bounded domain. Each domain was assigned a wave identification number (ID) but, if the COMs of two domains were less than a distance, Δ_{dom} , apart, they were given the same ID. If a domain's COM was within Δ_{dom} proximity of a domain's COM in the previous time bin, it was given the same ID. In this way, wave COMs were tracked according to their ID. After visually inspecting the results of the wave tracking algorithm for a number of waves with different values of Δ_{dom} , we found that $\Delta_{\text{dom}} = 680\mu\text{m}$ worked well. We used this value to analyse all complex retinal wave patterns simulated in this thesis. This algorithm performed well in separating waves that eventually merged with, or split from, other waves. This allowed us to perform analysis on every wave that had a unique ID, and to remove concurrent waves from that analysis.

To analyse firing rate movies of isolated retinal waves, we developed the space-time cross-correlogram (sCCG). To obtain the sCCG, we first computed the autocorrelogram (ACG) for the firing rate movie, $M(x, y, t)$, in which all but the one wave being analysed were removed. The ACG is defined as

$$ACG(\Delta x, \Delta y, \Delta t) = \frac{A(\Delta x, \Delta y, \Delta t)}{T(\Delta x, \Delta y, \Delta t)\lambda_M}, \quad (2.7)$$

where A is the autocorrelation of M :

$$A(\Delta x, \Delta y, \Delta t) = \sum_{i=1}^{N_x} \sum_{j=1}^{N_y} \sum_{k=1}^{N_t} M(x_i, y_j, t_k) M(x_i + \Delta x, y_j + \Delta y, t_k + \Delta t). \quad (2.8)$$

λ_M is the mean of the squared firing rates in M :

$$\lambda_M = \frac{1}{N_x N_y N_t} \sum_{i=1}^{N_x} \sum_{j=1}^{N_y} \sum_{k=1}^{N_t} M(x_i, y_j, t_k)^2, \quad (2.9)$$

and T is a triangle function that accounts for the of overlap of finite data as M is shifted by $(\Delta x, \Delta y, \Delta t)$ in Eq. 2.8:

$$T(\Delta x, \Delta y, \Delta t) = (N_x - |\Delta x|)(N_y - |\Delta y|)(N_t - |\Delta t|). \quad (2.10)$$

Note that the temporal ACG, averaged over all units, exists along $ACG(0, 0, \Delta t)$ and the spatial ACG, averaged over all time bins, exists along $ACG(\Delta x, \Delta y, 0)$.

Our goal was to compute the mean temporal CCG for two input units as a function of the distance between them. Therefore, we transformed the ACG into cylindrical coordinates, where $\Delta x = \Delta r \cos(\theta)$ and $\Delta y = \Delta r \sin(\theta)$. The two spatial dimensions were then collapsed into one to obtain the sCCG by integrating over θ :

$$sCCG(\Delta r, \Delta t) = \int_0^{2\pi} ACG(\Delta r, \theta, \Delta t) d\theta. \quad (2.11)$$

In this form, the sCCG is symmetric in Δt , so we use only $sCCG(\Delta r, \Delta t > 0)$. Thus, the sCCG provides a measure of correlation strength between the firing rate activity of two units, separated by a distance in space and time of $(\Delta r, \Delta t)$.

Travelling waves are characterised by a ridge of strong correlations in the sCCG. The Δr and Δt coordinates along the crest of the ridge have an approximately

linear relationship that is expected to reflect the speed of the wave. To determine the $(\Delta r, \Delta t)$ coordinates along the ridge crest, we found the Δr coordinate of the highest peak in the sCCG for each time bin. However, correlations from within the wavefront itself disrupt this linear relationship in the vicinity of (0,0). We wished to avoid this region of the sCCG and focus only on the linear region of the ridge, where correlations due to wave propagation are not swamped by correlations from within the wavefront. To achieve this, we computed linear regressions on all possible contiguous portions of the set of sCCG peaks. Of the regressions that gave a R^2 statistic ≥ 0.9 , we chose the regression that covered the largest contiguous set of peaks. Not all waves produced ridges that satisfied $R^2 \geq 0.9$. These waves were deemed non-travelling.

The wave speed was determined by the inverse of the gradient for the chosen regression and the wave size corresponded to the largest Δr value that was used in the chosen regression. Here, the wave size refers to the distance over which strong correlations are propagated. Thus, for non-travelling waves, the wave speed and size could not be determined.

As an alternative measure of wave size for non-travelling waves, we measured the wave area. This was simply the number of input neurons that were recruited by a single wave, multiplied by the square of the input neuron spacing, $1156\mu\text{m}^2$.

Converting between visual angle and retinal distance

To set up a number of simulations throughout the thesis, we use data from experimental studies regarding receptive fields, or travelling wave speeds, for example, which require conversion between degrees of visual angle and units of distance along the retinal surface. The precise conversion relationship between visual angle and

retinal distance depends on the geometry of the eye, which differs for different animals. We therefore use a general conversion factor to provide a useful estimate. To compute this conversion, we use the model eyeball depicted in Fig. 2.1, which we consider to be a sphere that collects light over 180° of visual angle. We assume that all of the light is refracted onto the retina, which extends from the fovea (point F in the figure) to the ora serrata (point O), subtending an angle of 120° between the F, O and the centre of the eye, C (Sterratt et al., 2013). The retinal distance, R , and visual angle, A , measured in degrees, can be related by the following expression:

$$\begin{aligned} R &= 2rA \cdot \frac{2}{3} \cdot \frac{\pi}{180} \\ &= \frac{4\pi r}{540}A, \end{aligned} \tag{2.12}$$

where r is the radius of the eyeball and the fraction, $2/3$, accounts for the circumferential coverage of the eyeball by the retina. Using $\angle OCF = 120^\circ$, we estimated the following radii for the eyes of the following animals at different stages of development:

- Mouse, 1–7 days old: $r \approx 1.0$ mm (Table 1 in Bodenstein & Sidman)
- Mouse, adult: $r \approx 1.7$ mm (Table 1 in Bodenstein & Sidman)
- Ferret, eye opening: $r \approx 2.8$ mm (Fig. 8 in Henderson et al., 1988)
- Cat, birth: $r \approx 4$ – 5 mm (Fig. 3 in Thorn et al., 1975)
- Cat, 4 weeks old: $r \approx 5$ – 6 mm (Fig. 3 in Thorn et al., 1975)
- Cat, adult: $r \approx 10$ mm (Fig. 3 in Thorn et al., 1975)




Figure 2.1: **Schematic horizontal section of the eyeball, for converting between degrees of visual angle and retinal distance.** Light rays incoming from a visual field spanning 180° enter the eye through the pupil and fall on the retina (thick line), which spans the anterior side of the eye from the fovea to the ora serrata. P - pupil; C - centre of the eye; F - fovea; O - ora serrata; r - eyeball radius.

3. Results

Theoretical principles for spike-timing dependent plasticity of neural circuits driven by travelling wave activity patterns

Background

Travelling waves of electrical activity are a prominent feature of the developing brain (Leinekugel et al., 1998; Garaschuk et al., 2000; Conhaim et al., 2010; Watt et al., 2009; Meister et al., 1991). That they induce strong correlations between nearby neurons has prompted the hypothesis that they may act through mechanisms of plasticity that depend on correlated activity in a manner proposed by Hebb (1949), to instruct and refine the patterning of connections in neural circuits.

A candidate mechanism for Hebbian development is spike-timing dependent plasticity (STDP), whereby changes in synaptic strength depend upon the relative timing

of pre- and postsynaptic spikes (Markram et al., 1997; Zhang et al., 1998), within a narrow temporal window with a time constant that is typically tens of milliseconds. Broadly speaking, STDP rules can be categorised into those that depend on the temporal order of pre- and postsynaptic spikes, called asymmetric STDP, and those that do not, called symmetric STDP. This dependence can be sensitive to many factors, including the type of synapse at which it operates, the location of the synapse on the postsynaptic cell, the precise pattern of spikes arriving at the synapse, the depolarisation and conductance state of the postsynaptic neuron, and the presence or absence of neuromodulators (reviewed extensively in Caporale & Dan, 2008, and Feldman, 2012). Despite these additional complexities, valuable insights regarding the potential utility of STDP in the nervous system have been provided with the use of simplified phenomenological STDP models.

Theoretical examination and network simulations of STDP have shone light on its ability to engage correlated input patterns to build structure into neural circuits (Kempster et al., 1999; Song et al., 2000; Song & Abbott, 2001; Gjorgjieva et al., 2009; Gjorgjieva et al., 2011) and to create receptive fields that resemble those in the visual system (Clopath et al., 2010). Analyses in these models have treated spatial and temporal correlations as separable features when considering their interaction with STDP. However, the spatiotemporal correlations induced by travelling waves are space-time inseparable. A number of studies have presented neural network simulations in which space-time inseparable activity patterns, including travelling waves, interact with temporally asymmetric STDP rules. Such works have demonstrated the utility of this interaction in building into the network a sensitivity for space-time inseparable stimuli, including motion in visual environments (Buchs & Senn, 2002; Rao & Sejnowski, 2003; Wensch et al., 2005; Honda et al., 2011), or specific temporal sequences such as those found in bird song (Fiete et al. 2010), by way of creating an asymmetry in the strength of lateral connections between pairs

of neurons, or by establishing a spatial shift in the location of strong feedforward connections. Senn & Buchs (2003) addressed the latter of these phenomena in a theoretical framework, showing that the spatial profile of synaptic strength changes, due to a single travelling pulse in a feedforward circuit, reflects the shape of the STDP rule. However, their analyses were constrained to a very particular network structure, and focused on synaptic changes due to sinusoidal input patterns, for their specific application to the development of direction selectivity in primary visual cortex. Thus, a formal analysis of the interaction between more generalised space-time inseparable activity patterns and STDP, and their long term effect on the spatial structure of generic feedforward networks, is lacking.

Here, we derive a mathematical expression for the interaction of travelling wave activity patterns with STDP in a simple yet biologically plausible model of spiking neurons. We first show that this interaction maps the time dependence of STDP into synaptic strength changes over space. We then show how such interactions can build periodic synaptic connectivity patterns into feedforward circuits. Our theoretical results hold for a broad class of STDP rule. Using simulations of spiking neurons, we verify our analytical results using two widely reported forms of STDP, and demonstrate the robustness of the pattern forming process to several likely sources of physiological noise. Because of the constraints imposed in the model of Senn & Buchs (2003), the authors find that changes in the connectivity pattern depend very little on the STDP time constant and the speed of sinusoidal activity patterns. However, by considering a more general case, our results reveal that the wave speed and the STDP time constant exert a strong influence over the resulting connectivity pattern. Moreover, the mechanism of pattern formation we describe here is analogous to a Turing instability in reaction-diffusion systems (Turing, 1952), whereby the initially homogeneous distribution of synaptic strengths becomes unstable, allowing the fastest growing eigenmode to dominate the resulting pattern.

Methods

Conceptual framework

To understand how correlated activity caused by travelling waves could influence synaptic strengths via a STDP mechanism, we begin by considering a reduced model (Fig. 3.1A) consisting of a 1D layer of presynaptic input cells, all of which connect via excitatory synapses onto a single, postsynaptic output cell. In following chapters, we extend this model to a 2D input layer. Consider the case of a single wave traversing the 1D input population (Fig 3.1A, red unit with rightward arrow). When the wave recruits each input neuron, that neuron discharges a burst of spikes. The set of spikes for all input neurons leads to a simple diagonal band structure in space and time (Fig. 3.1B). Synapses from the input layer to the output neuron are modified by a STDP rule, $K(\Delta t)$, which changes synaptic strengths according to the temporal difference, Δt , between input and output spike times (see General Methods, Eq. 2.3 and 2.4). For the purpose of this thesis, we use two distinct STDP rules: one that is asymmetric (Zhang et al., 1998; Feldman, 2000; Song et al., 2000; Fig. 3.1C, top) and one that is symmetric (Nishiyama et al., 2000; Wittenberg & Wang, 2006, Fig. 3.1C, bottom) in Δt . In the following, we use the symmetric STDP rule (a difference of Gaussians, see General Methods) to illustrate the sequence of events that lead to spatially structured changes in synaptic strengths for which we later derive more general, analytical solutions.

Fig. 3.1D (top) shows a snapshot when the travelling wave is in the middle of the input layer. At this instant, an input neuron is recruited by the wavefront and undergoes a burst of spikes, which elicit excitatory postsynaptic potentials (EPSPs) in the output neuron, thereby increasing the firing rate of the output neuron. Consequently, the increased output firing rate strengthens synapses from nearby in-




Figure 3.1: **Model framework and dynamics.** **A.** The model consists of a layer of input neurons, along which wavefronts of spiking activity propagate. Input neurons undergo a burst of spikes when the wavefront passes (red unit). Input spikes elicit excitatory postsynaptic potentials (EPSPs) in the output neuron (blue unit) that are scaled by the synaptic strength, w_{ij} . The output neuron generates spikes stochastically with a firing rate linearly proportional to its summed EPSP. **B.** Example spike raster from the simulation, showing one wave travelling past input neurons 1–10. **C.** Synaptic strengths are modified by either an asymmetric (top) or symmetric (bottom) STDP rule. **D.** Schematic for the mapping of a timing-dependent plasticity rule into a location-dependent rule. Top: an input neuron (red) is recruited by a wave and increases the firing rate of the output neuron (blue). When a symmetric STDP rule is at play, output spikes at the current time point will cause the synaptic strength of nearby input neurons to increase ('+' symbol) because they become active shortly before/after. Similarly, the strength of synapses from input neurons further away will decrease ('-' symbol). Thus, travelling waves map time-dependent into location-dependent plasticity rules. Bottom: at a later time point, the wavefront has moved and synapses that were previously strengthened are now weakened.

put neurons that were only recently recruited, or are soon to be recruited, by the same wave, because the STDP rule specifies strengthening for short Δt . Synapses connecting more distant input neurons will be weakened because the longer input-output Δt will fall under the negative lobe of the STDP rule. In fact, due to the wave motion, the time axis of the STDP rule, Δt , is mapped onto the spatial axis of the input layer, Δx , as shown in Fig. 3.2A. As such, the STDP rule can be reinterpreted as a spike-location dependent rule for travelling input activity patterns. Fig. 3.1D (bottom) shows a snapshot at a later time, when synapses that were previously strengthened are now weakened because of the output neuron's response to the new input activity. Overall, the net effect of a single, passing wave on synaptic strengths is not intuitive when looking at snapshots. One might predict that all synapses will change by the same amount, an amount that is proportional to $\int_{-\infty}^{\infty} d\Delta t K(\Delta t)$ and the horizontal axis. In the following analytical derivation, we show that this is not the case. Moreover, we show that that wave-induced STDP can lead to pattern formation in the network.

Travelling waves in the input layer

Unless otherwise stated, waves were made to travel with speeds of 4 mm/s. This falls within a broad range of wave speeds measured in different areas of the brain, from retinal waves at ~ 0.2 mm/s (Feller et al., 1997), to 7 mm/s waves in the hippocampus (Leinekugel et al., 1998) and 16 mm/s waves in the cortex (Conhaim et al., 2010). In some simulations, input neurons were given background firing rates. In these cases, the background firing rate only applied to neurons that were not at that time being recruited by a wave. Neurons that were recruited by a wave always had firing rates of 50 Hz and, in some simulations, 100 Hz. Otherwise, when an input neuron was not being recruited by a wave, it was silent. Consecutive waves were always made to travel in alternating directions so as to prevent drift of the connectivity patterns that developed. The exception to this was in simulations in which multiple

waves were present. This ensured that different waves did not cross each others' paths.

Theoretical results

Deriving synaptic strength dynamics

Here, we wish to derive a description for the dynamics of synaptic strengths that are driven by pairs of input and output spikes acting through a STDP rule, as a result of travelling wave activity patterns traversing the input layer. Within our framework, input and output spike trains, $S(t)$, are generated from a time dependent firing rate via a stochastic process. The stochastic arrival of spikes result in stochastic changes to the synaptic strengths, thus posing a considerable challenge when seeking a description for the pattern of synaptic strengths that evolves slowly over long periods of time, during which many travelling waves occur. It is therefore useful to separate the slow dynamics from the fast, stochastic dynamics under the assumption that, during a finite period of time, T , individual changes in synaptic strengths are negligible, but accumulate slowly over multiple T periods as a result of the time-averaged input and output activity. By approximating synaptic strengths as being constant during the period T , Kempster et al. (1999) showed that changes in synaptic strength over this period could be described by the inner product of the STDP rule, $K(\Delta t)$, and the cross-correlation function, $C_{i,j}(\Delta t, t)$, between the spike trains of input neuron i and output neuron j :

$$\frac{w_{ij}(t+T) - w_{ij}(t)}{T} = \frac{\Delta w_{ij}(t)}{T} = \eta \int_t^{t+T} d\Delta t K(\Delta t) C_{ij}(\Delta t, t), \quad (3.1)$$

where η is a small, positive constant that sets the required slow rate of change in synaptic strengths. The cross-correlation is given by:

$$C_{i,j}(\Delta t, t) = \int_t^{t+T} dt' S_i(t' + \Delta t) S_j(t'), \quad (3.2)$$

where $S_{i(j)}(t)$ are the input (output) spike trains. Given that we will deal with discrete waves that pass one-by-one across the input layer, it is convenient to relate the time scale, T , to the passage time of a single wave. As such, Eq. 3.1 implements the above approximation by essentially averaging over the small stochastic fluctuations, thus providing only the mean drift in $w_{i,j}(t)$ as a result of a single wave. This enables $C_{i,j}(\Delta t, t)$ to be calculated easily, as $S_j(t)$ depends on the synaptic strengths, $\sum_i w_{i,j}(t)$. Further assumptions are required to uphold the validity of Eq. 3.1. First, in order to ensure that multiple waves do not mutually influence changes in synaptic strength, T must be much longer than the temporal width of the STDP rule, which we denote as T_K and is typically tens of milliseconds. Because we are effectively considering $\Delta w_{i,j}(t)$ for a wave in isolation, with $T \gg T_K$, we can extend the integral limits in Eqs. 3.1 and 3.2 to $\pm\infty$. Second, the input and output firing rates must be sufficiently high in order for $C_{i,j}(\Delta t, t)$ to accurately portray wave-induced correlations. Alternatively, η must be sufficiently small for wave-induced correlations to be recovered over several waves with low firing rates. In either case, we can reinterpret $S_{i(j)}(t)$ as the time dependent firing rates corresponding to a single wave, rather than spike trains. Furthermore, with $w_{i,j}(t)$ effectively constant during a single wave, and because $\Delta w_{i,j}(t)$ is small, we approximate $w_{i,j}(t)$ with $w_{i,j}(T)$ and $C_{i,j}(\Delta t, t)$ with $C_{i,j}(\Delta t, T)$, and therefore approximate the left-hand side of Eq. 3.1 with $\partial w_{i,j}(T)/\partial T = \partial_T w_{i,j}(T)$ (Kempster et al., 1999).

In our simplest model, the input population is organised along a 1D chain, and we consider only a single output neuron, so we replace $w_{i,j}(T)$ with $w(x, T)$ and

$C_{i,j}(\Delta t, T)$ with $C(x, \Delta t, T)$. We can then write:

$$\partial_T w(x, T) = \eta \int_{-\infty}^{\infty} d\Delta t K(\Delta t) C(x, \Delta t, T), \quad (3.3)$$

where $C(x, \Delta t, T) = \int_{-\infty}^{\infty} dt S_{\text{in}}(x, t + \Delta t) S_{\text{out}}(t)$, with $S_{\text{in}}(x, t)$ and $S_{\text{out}}(t)$ the input and output firing rates, respectively. As a final matter of notation, we will hereafter refrain from explicitly writing the dependence of w and C on T , for brevity.

We make two further assumptions to simplify the analytical derivation, then relax these for a more general case: *i*) the input firing rate at the wavefront can be described as a short, travelling pulse using a Dirac delta function: $S_{\text{in}}(x, t) = \delta(x - vt)$, where $\delta(y) = 1$ if $y = 0$ and is zero otherwise, and v is the wave speed; *ii*) the output neuron's response to its input is instantaneous: $S_{\text{out}}(t) = \int_{-\infty}^{\infty} dx S_{\text{in}}(x, t) w(x) = w(vt)$. Using these forms for $S_{(\text{in})\text{out}}$, $C(x, \Delta t) = w(x - v\Delta t)$ and we can write Eq. 3.3 as

$$\begin{aligned} \partial_T w(x) &= \eta \int_{-\infty}^{\infty} d\Delta t K(\Delta t) w(x - v\Delta t) \\ &= \frac{\eta}{v} \int_{-\infty}^{\infty} d\Delta x K(\Delta x/v) w(x - \Delta x) \\ &= \eta K_v(x) * w(x), \end{aligned} \quad (3.4)$$

where $K_v(x)$ has been introduced as a rescaled copy of K , $K_v(x) = v^{-1}K(x/v)$. There are two key features to Eq. 3.4. Firstly, as illustrated above, the STDP rule can be reinterpreted as a function of space, as a result of the wavefront's velocity, which maintains a strict relationship between space and time. Secondly, the wave's effect on the synaptic dynamics is described by a convolution of the STDP rule with the synaptic strengths. In deriving a solution to Eq. 3.4, we will demonstrate in the following sections how convolution plays an important role in the type of connectivity patterns that $w(x)$ acquires. First, we consider a more realistic model that includes presynaptic bursts and EPSPs. These additional features add just a

simple modification to Eq. 3.4, which becomes

$$\partial_T w(x) = \eta K_v(x) * \alpha(-x/v) * \alpha(x/v) * \epsilon(x/v) * w(x), \quad (3.5)$$

where $\alpha(t)$ is the input firing rate during a burst and $\epsilon(t)$ is the EPSP. The full derivation for Eq. 3.5 is provided in the Appendix of this chapter. Note that the firing rate of the output neuron (set by r in Eq. 2.2) acts simply as a coefficient of ϵ in Eq. 3.5 and hence plays a similar role to η by varying the rate at which synaptic strengths are modified. Amalgamating all terms in Eq. 3.5, except for $w(x)$, we have

$$\partial_T w(x) = \eta \kappa(x) * w(x), \quad (3.6)$$

where $\kappa(x)$ is the spatial interpretation of the effective STDP rule:

$$\kappa(x) = K_v(x) * \alpha(-x/v) * \alpha(x/v) * \epsilon(x/v). \quad (3.7)$$

A solution for $w(x)$ as $T \rightarrow \infty$ is more easily found in the frequency domain by taking the Fourier transform, such that convolution becomes multiplication:

$$\partial_T \tilde{w}(k) = \eta \tilde{\kappa}(k) \tilde{w}(k), \quad (3.8)$$

where ‘ $\sim$ ’ denotes the Fourier transform and k is the spatial frequency. This is a first order differential equation with a well known solution. Using w_0 to describe the initial synaptic strengths at time $T = 0$,

$$\tilde{w}(k) = \tilde{w}_0(k) e^{\eta \tilde{\kappa}(k) T}. \quad (3.9)$$

By writing $\tilde{\kappa}(k) = \lambda(k) + i\phi(k)$ to include the real and imaginary components of $\tilde{\kappa}(k)$, we have

$$\tilde{w}(k) = \tilde{w}_0(k)e^{\eta\lambda(k)T}e^{i\eta\phi(k)T}, \quad (3.10)$$

we see that the stability of $\tilde{w}(k)$ is determined only by the real component, $\lambda(k)$, whereas the imaginary component gives rise to a phase shift in $w(x)$. For any k such that $\lambda(k) < 0$, $\tilde{w}(k)$ will be stable and decay to zero, whereas for any k such that $\lambda(k) > 0$, $\tilde{w}(k)$ will become unstable and grow exponentially.

For the STDP rules used in this study (Fig. 3.1C), $\tilde{\kappa}$ has a positive valued maximum at k^* (Fig. 3.2B). Therefore, if the synaptic strengths are initially random, such that the expected form of $\tilde{w}_0(k)$ is flat, then, as T increases, $\tilde{w}(k)$ will grow faster at $k = k^*$ than at any other k and will asymptotically approximate a delta function, $\delta(k^*)$. If only a single spatial frequency dominates $\tilde{w}(k)$, $w(x)$ will be sinusoidal. Thus, our derivation predicts that, for STDP rules like those reported experimentally, the synaptic strengths will develop a connectivity pattern that is periodic in space with a period of $1/k^*$, under the influence of travelling waves. We refer to k^* as the dominant spatial frequency of the connectivity pattern.

A similar model of pattern formation describes the development of ocular dominance columns in primary visual cortex (Swindale, 1980), in which it is proposed that short range excitatory and long range inhibitory lateral interactions give rise to the spatially periodic dominance of eye specific afferents across primary visual cortex. In our model, it is the mapping of the STDP rule onto space that provides the lateral interactions necessary for pattern formation. More generally, the mechanism of pattern formation we have presented is analogous to a Turing instability in reaction-diffusion systems (Turing, 1952), whereby the initially homogeneous

distribution of synaptic strengths becomes unstable, allowing the fastest growing eigenmode to dominate the resulting pattern. In systems composed of a diffusible activator and an inhibitor, Turing instabilities can only occur when diffusion of the inhibitor is greater than diffusion of the activator (Meinhardt & Gierer, 2000). Thus, STDP rules with narrow windows for potentiation and wide windows for depression are good candidates for pattern formation in neural circuits that support travelling wave activity patterns.

Limitations to the induction of periodic patterns by waves and STDP

The contraction of time scales such that we can treat T as the continuous variable places restrictions on the properties of travelling waves for which our analytical results are valid. Given the end goal of obtaining a periodic connectivity pattern, we explore the limits these restrictions impose on network dynamics such as the wave frequency, and the duration of input bursts, in the last three sections of the Results, and show that pattern formation can be achieved over a wide range of biologically realistic waves.

Simulation results

Travelling waves induce spatial pattern formation in a spiking neuron simulation

To test the predictions made by the analytical derivation above, we examined whether a periodic connectivity pattern would develop as a result of travelling waves and STDP in a simulated network of linear Poisson, spiking neurons. The simulation captures the architecture and function of the network illustrated in Fig. 3.1A and consists of a 1D layer of input neurons that all synapse onto a single output neuron. During simulations, synaptic strengths are modified by a STDP rule, ac-

cording to Eq. 2.3 or 2.4 in General Methods. Spiking activity is generated by a travelling wavefront that moves back and forth along the input layer, eliciting EPSPs in the output neuron, causing it to spike. We varied two parameters in the simulation: *i*) the temporal window of $K(\Delta t)$, which is typically controlled with the decay time constant for synaptic strengthening, τ_+ , and *ii*) the wave speed, v . These parameters modulate the shape of $\kappa(x)$ and, therefore, the spatial frequency of the predicted periodic connectivity pattern. We ran two sets of simulations, one using the asymmetric STDP rule (Zhang et al., 1998; Feldman, 2000; Song et al., 2000) and the other using a symmetric rule (Nishiyama et al., 2000; Wittenberg & Wang, 2006).

Fig. 3.2C shows that the synaptic strengths in the simulation indeed developed a periodic connectivity pattern in the presence of travelling waves for both the asymmetric and symmetric rules, as predicted by Eq. 3.9. The periodic pattern was robustly produced over a range of STDP time constants ($\tau_+ = 20, 30, 40, 50, 60$ and 70 ms), keeping the wave speed constant ($v = 3$ mm/s), and for a range of wave speeds ($v = 1, 2, 3, 4, 5$, and 6 mm/s), keeping τ_+ constant ($\tau_+ = 20$ ms). Specifically, increasing τ_+ or v caused the period of the connectivity pattern to increase as expected (Fig. 3.2C). To test the accuracy with which Eq. 3.9 predicts the resulting spatial frequency over this range of parameters, we computed the power spectrum of the steady state synaptic strengths (mean subtracted) to which a Gaussian curve was fit. The spatial frequency of the connectivity pattern was taken to be the centre of the Gaussian fit. We repeated simulations sixteen times for each set of parameters, using a different seed for the random number generators. In Fig. 3.2D, we compare the measured k^* (blue circles) with the predicted k^* (grey line) as a function of τ_+ (top panels) and v (bottom panels), where the predicted values were found by numerically computing $\tilde{\kappa}(k)$. We computed the coefficient of determination, R^2 , between the predicted values of $\log(k^*)$ and those measured




Figure 3.2: **Travelling waves build periodic connectivity patterns into the synaptic strengths through STDP.** **A.** Example of a STDP rule (top panel, $\tau_+ = 20$ ms) that is mapped into location-dependent plasticity rule (bottom panel) by a wave travelling at 3 mm/s. **B.** $\text{Re}[\tilde{\kappa}(k)]$ for the STDP rule in **A**. The spatial frequency where $\tilde{\kappa}(k)$ has a global maximum is k^* , which we call the dominant spatial frequency. Inset: $\tilde{\kappa}(k)$ for a symmetric STDP rule ($\tau_+ = 20$ ms, $v = 3$ mm/s). **C.** Evolution of the synaptic strengths, initialised with strength 0.5, from 500 input neurons to a single output neuron, under the influence of an asymmetric (top row) and symmetric (bottom row) STDP rule. Waves sculpt a periodic connectivity pattern into the synaptic strengths with a spatial frequency equal to k^* (see **B** and Eq. 3.5). Larger values for τ_+ (controls STDP rule width) or v (wave speed) produce lower spatial frequencies in the periodic pattern. Values for τ_+ and v given between the top and bottom rows. **D.** Log-log plots for the spatial frequency of the steady state connectivity pattern as a function of τ_+ (top) and v (bottom) for both symmetric (left) and asymmetric (right) STDP rules. Grey line: predicted dominant spatial frequency, k^* ; blue: mean spatial frequency $\pm$ SEM of the final connectivity pattern in the spiking simulation; purple: mean spatial frequency $\pm$ SEM of $w(x)$ in the numerical solution to Eq. 3.6.

from the simulated data as an indicator of accuracy. For each panel in Fig. 3.2D, $R^2 > 0.85$. Thus, the assumptions made to derive Eq. 3.9 appear to maintain a veritable description of the noisy dynamics in the simulation over the range of parameters tested. In addition to the simulations, we verified that numerical solutions to Eq. 3.6 (which included nonlinearities such as hard bounds to $w(x)$, see General Methods) also produced periodic connectivity patterns with spatial frequencies that matched predictions (magenta circles, Fig. 3.2D, $R^2 > 0.92$). In a later section and in Chapter 4, we use numerical solutions to more completely sample the broad parameter space in which experimentally observed travelling waves exist.

The development of periodic patterns in synaptic connectivity could have wide applications throughout the nervous system of many species, particularly because of the ubiquitous findings of the two major phenomena: waves and STDP. However, it is first necessary to determine how robust this process is 1) across a range of spatial and temporal scales, and 2) in the presence of noise. We explore these issues in the following sections.

Pattern formation over a wide range of spatial and temporal scales

So as to address the applicability of the above theoretical results to biological wave-mediating circuits, we calculated the landscape of spatial frequencies acquired by the connectivity patterns as τ_+ and v are varied over a broad range by numerically computing the dominant spatial frequency, k^* , which maximises the real component of $\tilde{\kappa}(k)$, for the two STDP rules used above. We deliberately extend these calculations to STDP time constants that are long enough to make them applicable to the relevant scales for retinal waves and to spike bursts that are often found to accompany spontaneous waves at early developmental stages.

Fig. 3.3 shows the contour maps of the k^* landscapes for the asymmetric and sym-

metric STDP rules examined throughout this thesis. Given the simple exponential functions used to describe these STDP rules (Eq. 2.3 and 2.4 in General Methods), it is perhaps not surprising that the spatial frequencies they impose (Figs. 3.3A and B) depend on τ_+ and v in a remarkably similar way. In particular, for a large region of parameter space, scaling either v or τ_+ by a factor, f , simply scales the dominant spatial frequency, k^* , by $1/f$. We illustrate this scaling feature using triangles with two equal sides aligned to the axes, as multiplication in linear space is equivalent to addition in logarithmic space. Starting at one iso-frequency contour (1 cycle/mm, bottom left vertex of white triangle in Fig. 3.3A), a constant step along the v -axis moves k^* to the same iso-frequency contour (0.04 cycles/mm) as does an equal step along the τ_+ -axis. This relationship does not continue when τ_+ is shorter than the burst duration of input neurons (in this case, 0.1 s; green triangle in Fig. 3.3A). Thus, the burst duration has a strong influence on k^* at times scales longer τ_+ . Burst durations much longer than τ_+ also challenge one assumption made to derive our theoretical results above: that individual synaptic strengths change negligibly during a wave in comparison to changes over successive waves. Long burst durations of 1–3 s are a common feature of retinal waves (Stafford et al., 2009; Maccione et al., 2014). Addressing the impact of long burst durations is the focus of the next section in this chapter, and an examination of different STDP rules that may be appropriate for waves with long burst durations is given in Chapter 5.

Travelling waves in different brain areas do not all have the same features, yet we may ask similar questions about each of them regarding the impact they would make on the development of efferents from those areas, according to our theoretical results. Waves in the ventral neonatal mouse cortex (Conhaim et al., 2010) and hippocampus (Leinekugel et al., 1998) travel with speeds of ~ 7 mm/s, which, when coupled with asymmetric STDP rules recorded in mouse cortex (Feldman, 2000; Froemke et al., 2005), can be positioned in Fig. 3.3A at the red dot. We would therefore predict

that, as a results of these waves, the spatial wavelength of the resulting connectivity pattern with postsynaptic targets would be ~ 1.8 mm. This is approximately equal to the stretch of cortex along which ventral cortical waves propagate, and approximately half the distance travelled by hippocampal waves. Comparison of the connectivity pattern’s wavelength to the distance travelled by the wave is equivalent to comparing the passage time of the wave, T , to the time scale of the STDP rule, τ_+ . In order for our analytical results to hold, we had to assume $T \gg \tau_+$. In Chapter 5, we examine how the distance travelled by waves affects the connectivity pattern in the context of retinal waves, which propagate over a range of distances. In a later section in this chapter, we explore a related question regarding the influence on the connectivity pattern of multiple waves occurring within the period, T .

Retinal waves that occur during the first postnatal week in mice are typically observed with speeds of 0.1–0.3 mm/s (Stafford et al., 2009; Maccione et al., 2014), and have been implicated with the development of receptive field sizes in the mouse superior colliculus (SC) (Chandrasekaran et al., 2005). Chandrasekaran et al. reported receptive field sizes of mature, SC neurons in wild type mice to be $\sim 16^\circ$ of visual angle. A receptive field can be interpreted as half a cycle in a periodic connectivity pattern, an assumption we make in Chapter 4 to examine the application of our theoretical results to receptive field refinement. Under this assumption, and by converting wave speeds from units of mm/s to units of $^\circ$ /s (6–18 $^\circ$ /s in mice; see Eq. 2.12 in General Methods for conversion formula), STDP rules would need to operate over time scales ranging 0.3–1.0 s (solid ellipse in Fig. 3.3C) in order to establish SC receptive fields of this size. We examine the plausibility of STDP rules with long time scales, and other candidate STDP rules, in Chapter 5.

In the primary visual cortex, simple cells exhibit a periodic spatial structure (Hubel & Wiesel, 1962) and can be characterised as responding to a preferred spatial fre-




Figure 3.3: **Dependence of the predicted spatial frequency on wave speed and STDP rule in relation to biological examples of travelling wave.** In all plots, solid contours denote spatial frequencies equal to 10 raised to integer exponents. **A.** Landscape of the predicted spatial frequency for an asymmetric STDP rule as a function of v and τ_+ . Red point: wave speeds recorded in ventral cortex and hippocampus of neonatal mice. White triangle: when τ_+ is longer than the burst duration, scaling τ_+ or v by the same amount has the same effect on the spatial frequency. Green triangle: when τ_+ is shorter than the burst duration, the latter has a substantial influence on the spatial frequency. **B.** Landscape of predicted spatial frequencies for a symmetric STDP rule. **C.** Same as **A**, with units of space in degrees for application to the visual system. Solid red circle: retinal wave speeds and mature receptive field sizes in the mouse superior colliculus, which correspond with the encircled spatial frequencies, require that STDP rules have time scales of up to 1 s. Dashed red circle: simple cells in ferret visual cortex are most responsive to stimuli with the encircled spatial frequencies; these spatial frequencies can be obtained by waves with speeds that match typical speeds of visual motion and STDP time scales that are typical of the cortex.

quency (Movshon et al., 1978). A significant fraction of simple cells also express a preference for motion of visual stimuli in a particular direction (Hubel & Wiesel, 1962), but only if the animal experiences visual motion early in development (Blake-more & Sluyters, 1975; Li et al., 2006). Leading circuit models for simple cells posit that direction selectivity can be built from non-direction selective inputs that have both spatially offset receptive fields and different response latencies (Adelson & Bergen, 1985; De Valois et al., 2000; Saul et al., 2005), but how this spatial organisation is established during development is not yet clear. Ferrets are a choice model animal in which to study the development of direction selectivity (Li et al., 2006 ; Li et al., 2008; van Hooser et al., 2012). Ferret simple cells exhibit preferred spatial frequencies over a wide range, 0.04–0.6 cycles/° (Smyth et al., 2003). In Chapter 6, we extend our theoretical results to address the development of direction selective simple cell receptive fields, demonstrate how a broad range of stimulus velocities can give rise to receptive fields with spatial frequency tuning and direction selective responses, and propose that STDP rules with time scales of ~ 20 ms may be best suited for this task (dashed ellipse in Fig. 3.3C).

Impact of burst duration on the separation of time scales

Travelling waves elicited in the early stages of development are often characterised by strong and long lasting depolarisations, upon which bursts of action potentials are generated, lasting up to several seconds. Being much longer than the decay time of most STDP rules, these types of activity challenge the assumption that the time scale for changes in synaptic strength during a single wave is much slower than the time scale for changes over successive waves. In contrast, it is possible for sensory driven activity to be very precise, lasting tens of milliseconds, during which very few spikes can be generated. In this limit, our assumption that the time averaged firing rate is a sufficient approximation for the network’s spiking activity is also challenged. We therefore explored the robustness of our analytical results to these

contrasting scenarios using simulations in which the burst duration of input neurons was a free parameter.

Fig. 3.4A shows examples of the time evolution of the connectivity pattern for different burst durations. For bursts lasting 30-50 ms, the connectivity pattern had only a very weak periodic structure, with power spread over a broad range of spatial frequencies (Fig. 3.4B (light blue), C). Input spikes were very sparse, spiking only an average of 1.5 times per neuron per wave with a firing rate of 50 Hz. This challenges the assumption, made in the analytical derivation above, that the input and output firing rates are sufficiently high for $C(x, \Delta T)$ to be a good estimate of wave induced correlations. Burst durations lasting up to a few hundreds of milliseconds fostered stronger correlations, thus enabling the formation of the periodic connectivity pattern (top-middle panels). However, in simulations with burst durations of 1 s or longer, the time scale of a burst became comparable to the slower time scale of plasticity. As the burst duration increases further, activity in the inputs began to approximate simulations in Song et al. (2000), for which input spike trains were uncorrelated and thus there was no temporal or spatial structure within the input activity. Consequently, the connectivity pattern in our simulations developed little spatial structure for long burst durations.

In Fig. 3.4B, we plot the mean power spectra for the connectivity patterns that emerged in these simulations (short to long burst durations are coloured from light blue to black), which show a clear trend of decreasing spatial frequencies with increasing burst duration. To summarise the robustness of the connectivity pattern under different conditions, we denote the power spectrum $P(k)$ and measure the normalised peak power, $P(k^*)/P(k_0^*)$, i.e. the power at the dominant spatial frequency, k^* , normalised by the power at the dominant spatial frequency in a set of reference simulations, k_0^* . Here, we use simulations when the burst duration was 0.1 s as a




Figure 3.4: **Robustness of the periodic connectivity pattern to different input burst durations.** **A.** Examples of synaptic strength evolution over time. Numbers along the top indicate the input burst duration in seconds. Top row: input firing rate of 50 Hz, 20000 simulated seconds. Bottom row: input firing rate of 200 Hz, 15000 simulated seconds. **B.** Mean power spectra of the final connectivity patterns in simulations using input firing rates of 50 Hz, normalised by the peak power in each spectrum. Light blue to black corresponds to burst durations of 0.03 s, 0.05 s, 0.3 s, 0.3 s, 0.5 s, 1 s, 3 s and 5 s. **C.** Robustness of the periodic connectivity pattern to different burst durations, using $P(k^*)/P(k_0^*)$, where $P(k_0^*)$ was taken from simulations with a burst duration of 0.1 s. Black: input firing rate of 50 Hz. Grey: input firing rate of 200 Hz. *a*, *b* and *c* correspond to the labelled peaks in **B**.

reference. The robustness measure is plotted in Fig. 3.4C for burst durations from 30 ms to 5 s (black circles).

To test whether the sparsity of input spikes was the cause for poor robustness at short burst durations, we repeated the simulations, using a input firing rate of 200 Hz. Bursts lasting 30 ms now generated 6 spikes on average. Increasing the firing rate substantially improved robustness for burst durations up to 0.5 s, but bursts > 1 s in duration still failed to produce a periodic connectivity pattern. This suggests that if waves are to build regular patterned feedforward circuits through long depolarisations and bursts, as with retinal waves, then it can not be simply done via the fast STDP rules so far reported in the literature. In Chapter 5, we investigate rules that are able to operate over longer time scales and address their interaction with retinal waves.

The uncorrelated spiking within a burst is a source of noise that makes a substantial contribution to the spatiotemporal structure of wave-related correlations when the burst duration is long. Other sources of noise must also be considered. The noisy mechanisms of wave generation may initiate waves in quick succession, and spontaneous spikes that occur outside of a wave can further disrupt wave-related correlations. The extent to which our analytical results can be applied to biological systems therefore depends on their robustness to these additional contributions to the correlations between input neurons, $C(x, \Delta t)$. In the following sections, we seek to understand how a multiplicity of waves and spiking noise may impact upon the development of periodic connectivity patterns.

Sensitivity of periodic connectivity patterns to noise

Thus far, we have tested theoretical predictions using very idealised wavefronts. In reality, both spontaneous and sensory driven waves can be highly variable, with

a spectrum of wave speeds and background firing rates that are independent of waves. Here, we use simulations to establish the sensitivity of pattern formation to these sources of noise. To implement noise in the wave speed, we conducted a new set of simulations in which we assigned a new speed to each wave from a normal distribution. The distribution was centred on 4 mm/s with a variance of σ_v^2 , and we applied a minimum speed of 0.05 mm/s so that waves would not take too long to traverse the input layer. Examples of the connectivity pattern as it evolved in these simulations are provided in Fig. 3.5A and, in Fig. 3.5C, we plot the mean power spectra of the final connectivity patterns for different σ_v in the range 0-4 mm/s. As σ_v increases, the connectivity pattern tends towards lower spatial frequencies (cyan through to black). This is due to the fact that faster wave speeds recruit more input neurons per unit time. The number of spike pairs contributing to the update of a single synapse will therefore increase as wave speed increases, emphasising the influence of faster waves, which build in lower spatial frequencies. In Fig. 3.5, we plot $P(k^*)/P(k_0^*)$ as a measure of robustness of the connectivity pattern to variation in wave speed, where $P(k_0^*)$ is measured for the case when $\sigma_v = 0$. Despite the variation in wave speed, which corresponds to there being a spectrum of spatial frequencies impressed on the network, the connectivity retains a robust periodic structure (Fig. 3.5E).

We next tested the robustness of periodic connectivity formation in the presence of background spiking noise, which diminishes the relative contribution of wave-activity to $C(x, \Delta t)$ in Eq. 3.3. To implement background noise, input neurons in the simulation were given a firing rate of 50 Hz when recruited by a wave and a tonic firing rate, R_0 , otherwise. Examples of the connectivity patterns as they evolved are shown in Fig. 3.5B and mean power spectra of the final connectivity patterns are drawn in Fig. 3.5D. We use $P(k^*)/P(k_0^*)$ as a measure of robustness, where k_0^* is now the dominant spatial frequency when $R_0 = 0$ Hz. The connectivity pattern remains reasonably robust to background spiking for $R_0 < 2$ Hz (Fig. 3.5F, black




Figure 3.5: **Robustness of periodic connectivity pattern to variation in wave speed and uncorrelated input spikes.** **A & B.** Examples of synaptic strength evolution over time for simulations in which: **A)** wave speeds were drawn from a normal distribution with mean 4 mm/s and variance σ_v^2 ; **B)** a background firing rate of R_0 was assigned to input neurons. The firing rate during a wave was 50 Hz (top row) or 100 Hz (bottom row). **C & D.** Power spectra for the connectivity patterns, normalised by the maximum in each spectrum, for simulations in which: **C)** wave speeds were drawn from a normal distribution. Light blue to black corresponds to increasing σ_v from 0 mm/s to 4 mm/s, in steps of 0.4 mm/s; **D)** input neurons were assigned a background firing rate. Light blue to black corresponds to R_0 taking values of 0 Hz, 0.25 Hz, 0.5 Hz, 1 Hz, 2 Hz and 4 Hz, respectively. **E & F.** $P(k^*)/P(k_0^*)$ as a measure of robustness of the periodic connectivity pattern to: **E)** variation in wave speed, where $P(k_0^*)$ is for simulations with $\sigma_v = 0$; **F)** background spiking of input neurons, where $P(k_0^*)$ is for simulations with $R_0 = 0$. Black: 50 Hz firing rate during a wave. Grey: 100 Hz firing rate during a wave.

circles), with the dominant spatial frequency matching that predicted. Increasing the input firing rate during a wave burst to 100 Hz improved robustness of the periodic connectivity pattern, which, in this case, maintained a periodic structure for $R_0 < 4$ Hz (Fig. 3.5F, grey circles). Our results make the strong prediction that very low background firing rates are required for this type of pattern formation. During early, spontaneous waves, this is typically the case, and there is good reason to believe that they are low in early sensory systems as well (Olshausen & Field, 2004). We review evidence for this in the Discussion.

Sensitivity of periodic connectivity patterns to multiple, non-isolated waves

Spontaneous waves in early development are often temporally isolated with respect to the time scale of STDP, with inter-wave intervals (IWIs) lasting tens of seconds. The occurrence of sensory driven waves, however, is likely to match the temporal pattern of stimulus features to which a population of input neurons are tuned. During vision, for example, multiple luminance contours can be cast across the retina in quick succession when tracking a moving object. In the following, we investigate how the presence of multiple waves, which simultaneously drive spiking in the output neuron, impact the development of the connectivity pattern. To do this, we generated waves in the simulation with an exponentially distributed IWI and varied the mean IWI between simulations. At a speed of 4 mm/s, waves took 5 s to traverse the input layer. Thus, to ensure that multiple waves were present for the majority of the simulation, we set $IWI = 2.0$ s. In these simulations, waves were made to travel left to right only, such that they never crossed paths. Consequently, because of the asymmetry of $\kappa(x)$, each wave gave rise to a spatial shift in the connectivity pattern (Fig. 3.6A). Figs. 3.6A, C and E (black circles) show that, despite multiple waves driving the output cell, a periodic pattern emerged, and did so for values of IWI as short as 0.5 s. To measure robustness, we ran “reference” simulations in which waves were initiated with a fixed IWI of 10 s, rather than drawing IWIs

from a distribution. This ensured that only one wave was present at a time. The dominant spatial frequency that arose in the reference simulations was used for k_0^* in our measure of robustness. The connectivity pattern had little robustness when $\text{IWI} = 0.2\text{ s}$, and consisted of plateaus of weak synapses between occasional peaks of the periodic pattern (Fig. 3.6A). This latter case produced a low frequency peak in the power spectrum (Fig. 3.6C, black), reducing power at the dominant spectral peak (Fig. 3.6E, black circles in top panel).

Why is 0.2 s a limiting IWI for periodic pattern formation? Note that the wave travels less than one cycle of the periodic pattern ($1/k^* = 1.1\text{ mm}$) in 0.2 s , such that the vicinity of one wave, over which $\kappa(x)$ has an influence, often overlaps with that of another wave, i.e. $\text{IWI} < 1/vk^*$. Calculating k^* numerically, we find that $1/vk^* = 0.28\text{ s}$. We hypothesise that a lower bound to the IWI exists and is equal to $\text{IWI}_{\min} = 1/vk^*$. When IWIs are below this lower bound, the connectivity pattern loses its periodic structure. We ran additional sets of simulations to corroborate this. First, because the dominant spatial frequency, k^* , is inversely proportional to v , we expected that $\text{IWI}_{\min}$ is independent of the wave speed. Indeed, increasing the speed from 4 mm/s to 8 mm/s had little effect on the robustness of the periodic pattern to different IWIs (Fig. 3.6E, grey circles in top panel). $\text{IWI}_{\min}$ is, therefore, determined by the time scales of the components that compose $\kappa(x)$: the STDP rule, $K(\Delta t)$, the input firing rate pattern, $\alpha(t)$, and the EPSP, $\epsilon(t)$ (see Eq. 3.7). In another set of simulations, we therefore increased the spatial frequency of $\kappa(x)$ by decreasing the burst duration from 100 ms to 50 ms , such that $\text{IWI}_{\min} = 0.19\text{ s}$. Consequently, when $\text{IWI} = 0.2\text{ s}$, the connectivity pattern remained as robust as when $\text{IWI} = 2\text{ s}$ (Fig. 3.6E, black triangles in bottom panel), but degraded when $\text{IWI} = 0.1\text{ s}$. To ensure this effect was due to a change in $\kappa(x)$ and not the burst duration alone, we ran simulations with the 50 ms burst duration but a longer STDP time constant, $\tau_+ = 60\text{ ms}$. In this case, $\text{IWI}_{\min} = 0.5\text{ s}$ and the connectivity pattern

was no longer robust when $IWI = 0.2\text{s}$. These results provide evidence that the spatial scale of $\kappa(x)$, and therefore the temporal scales of $K(\Delta t)$, $\alpha(t)$ and $\epsilon(t)$, determine the maximum frequency with which waves can occur and still build periodic connectivity patterns. Indeed, this is a necessary assumption in our derivation of Eq. 3.6.

In experimental studies of visual development, a protocol often used to induce selectivity in visually responsive neurons is a high contrast, drifting grating that is presented to an animal's eye (for example, see Li et al., 2008). This type of stimulation likely activates multiple waves in the retina that track the high contrast edges of the stimulus. To emulate this paradigm, we ran simulations in which waves occur in a periodic fashion with a fixed IWI and found that the robustness of the periodic connectivity pattern was vastly improved over all IWIs. Fig. 3.6B shows examples of the time evolution of the connectivity pattern for different IWIs, using the same simulation parameters as for Fig. 3.6A, except that each wave was generated after a fixed IWI. Even when the IWI was 0.2s , the connectivity pattern remained robust (Fig. 3.6F). However, the peak in the power spectrum shifted to a higher frequency (Fig. 3.6D, black). These results appear to contradict our results from simulations using exponentially distributed IWIs, as well as the assumptions made for our theoretical derivation, but they can be easily accounted for by considering how $\kappa(x)$ is constructed in Eq. 3.7. This equation expresses $\kappa(x)$ as the convolution of $K(x/v)$, $\alpha(x/v)$ and $\epsilon(x/v)$, and can be solved by taking its Fourier transform, whereby convolution in the real domain is equivalent to multiplication in the frequency domain.

In Fig. 3.6G, we plot the two functions that have the greatest influence over pattern formation in the above simulations. In black is the Fourier transform of the STDP rule, $\text{FT}\{K(x/v)\}$. The remaining curves are the Fourier transforms of the input firing patterns (with the zero-frequency component removed), $\text{FT}\{\alpha(x/v)\}$,




Figure 3.6: Robustness of periodic connectivity pattern to inputs from multiple, simultaneous waves. **A & B.** Examples of synaptic strength evolution over time for simulations in which: **A)** waves are generated with a mean IWI drawn from an exponential distribution, with wave speeds of 4 mm/s; **B)** waves are generated regularly with a fixed IWI. **C & D.** Power spectra for the connectivity patterns, normalised by the maximum in each spectrum, for simulations in which: **C)** IWIs are drawn from an exponential distribution. Light blue to black corresponds to IWIs of 2.0 s, 1.5 s, 1.0 s, 0.5 s and 0.2 s, respectively; **D)** IWIs are fixed. Colour scheme the same as for **C**. **E & F.** $P(k^*)/P(k_0^*)$ as a measure of the connectivity patterns' robustness to different IWIs, where $P(k_0^*)$ is for simulations with fixed IWIs of 10 s. **E)** IWIs are exponentially distributed. Top panel: robustness to different IWIs has little dependence on wave speed. Bottom panel: robustness depends instead on the shape of $\kappa(x)$. Impact of burst duration is seen by comparing black markers in top and bottom panels. Impact of τ_+ is seen by comparing black and grey markers in bottom panel. **F)** IWIs are fixed. **G.** Real components in the Fourier transform of the STDP rule ($K(x/v)$, black) and the input firing rate function ($\alpha(x/v)$) for a single wave in isolation (grey) and multiple waves with IWIs of 1 s (blue) and 0.2 s (pink). **H.** Real components in the Fourier transform of $\kappa(x)$ for the same IWIs in **G**.

for three cases: a single wave in isolation (grey) and periodic waves with $IWI = 1.0$ s (blue) and $IWI = 0.2$ s (pink). Note that the Fourier transform of the EPSP has little effect in this calculation if its time constant is much shorter than that of the other functions. Multiplication of $FT\{\alpha(x/v)\}$ with $FT\{K(x/v)\}$ yields $FT\{\kappa(x)\}$, which is drawn for the three cases of $\alpha(x/v)$ in Fig. 3.6H. When $\alpha(x/v)$ is periodic, its Fourier transform is also periodic. This can produce strong peaks in $FT\{\kappa(x/v)\}$ (pink and blue) at spatial frequencies that may be higher than the peak in the case of isolated waves (pink). Large amplitude, narrow peaks in $FT\{\kappa(x)\}$ for periodic waves explain the dominance and robustness of a single spatial frequency in the connectivity pattern, and the positions of these peaks explains the shift of the power spectrum toward higher spatial frequencies when $IWI = 0.2$ s in Fig. 3.6D. However, when waves traverse the inputs on an irregular basis, $FT\{\alpha(x/v)\}$ assumes a noisy copy of the isolated wave case. As such, the calculation of $\kappa(x)$ and, therefore, the lower bound to the IWI, is effectively the same as for the isolated wave. Nevertheless, for both non-periodic and periodic waves, the IWI must still have a lower limit (or upper limit to the temporal frequency of the stimulus), which is set by the burst duration, such that multiple waves not lie on top of each other. In summary, the presence of multiple waves in biological circuits poses few problems for the applicability of our analytical results, and are tolerant to reasonably high wave frequencies that may be elicited by sensory stimulation.

Discussion

An instructive role for travelling waves in the developing brain has long been proposed. Such activity patterns induce local correlations in space and time, which can build structure into the connectivity of neural circuits through Hebbian mechanisms of plasticity. Here, we have built upon these ideas to show how travelling waves interact with STDP, and our analysis has led to the following major results:

1) common forms of STDP reported in the literature build periodic patterns into feedforward networks when driven by travelling waves; 2) the spatial frequency of the periodic pattern is inversely proportional to both the STDP time constant and wave speed; 3) waves that induce input bursts ≥ 1 s in duration fail to build periodic connectivity patterns through STDP rules with time scales of tens of milliseconds; 4) pattern formation by travelling waves and STDP is very robust to variations in wave speed and wave frequency, but is not tolerant of background spiking at rates exceeding 1–2 Hz.

Our theoretical results, expressed in Eqs. 3.6 and 3.9, are analogous to those derived by Swindale (1980) to model the development of ocular dominance columns in primary visual cortex. Swindale showed how short range excitatory and long range inhibitory lateral interactions could give rise to the spatially periodic dominance of eye specific afferents across primary visual cortex. In our model, it is the mapping of the STDP rule onto space that provides the lateral interactions necessary for pattern formation. This mechanism of pattern formation is analogous to a Turing instability in reaction-diffusion systems (Turing, 1952). A generally accepted requirement for Turing-like pattern formation is the presence of a diffusible activator and inhibitor, whereby diffusion of the inhibitor is greater than diffusion of the activator (Meinhardt & Giesler, 2000). Thus, STDP rules with narrow windows for synaptic potentiation and wide windows for synaptic depression are good candidates for pattern formation in neural circuits that support travelling wave activity patterns.

Using simulations of a spiking neural network, we examined the sensitivity of our analytical results to physiological conditions that are commonplace in the brain, and highlight several conditions under which periodic pattern formation cannot be obtained. In particular, waves that induce long bursts lose the ability to build periodic connectivity patterns, because correlations within single input bursts occlude wave

related correlations over the short time scale of the STDP rule. This is relevant to several types of spontaneous wave that have been observed in the developing brain (Leinekugel et al., 1998; Stafford et al., 2009; Conhaim et al., 2010), and suggests that the short time scale often reported for STDP (Caporale & Dan, 2008) is not suitable for building structured circuits from these waves. This idea has previously been expressed with regard to the role of retinal waves, which likely require plasticity rules that operate over long time scales if they are to play an instructive role in visual development (Butts & Rokhsar, 2001; Butts et al., 2007). Our work adds support to the prediction that STDP rules at retinofugal synapses must operate over time scales much longer than are often reported, if they are to engage retinal waves in a manner that instructs refinement of retinofugal connections.

Another important physiological state to consider during early development is that neural circuits are typically highly excitable and noisy, which can lead to moments when multiple waves traverse a population in quick succession and with variable wave speeds. This appears to have little impact on wave-induced pattern formation, so long as the mean IWI exceeds $1/vk^*$. Pattern formation is, however, sensitive to background noise in the form of non-wave related spikes. This sensitivity was dependent on the ratio of the background firing rate to the wave-induced firing rate. During stage 2 retinal waves in mice, retinal ganglion cells rarely spike outside of a burst. However, bursts that occur outside of a wave event have approximately the same firing rate as wave related bursts (Table 1 in Stafford et al., 2009). Stafford et al. found that wave related bursts comprise approximately 90% of all bursts, suggesting that background spiking noise would have little impact on wave-induced correlations at this stage of visual development. Low amplitude background noise appears to continue throughout visual experience in the retina, whereby retinal ganglion cells exhibit reliable response motifs to the same pattern of visual stimulation (Pitkow & Meister, 2012). This is especially the case for population responses in

the retina to moving visual contours (Frechette et al., 2005). Pitkow & Meister (2012) argue that the response behaviour of retinal ganglion cells is evidence for a sparse coding scheme. That sparse coding provides a successful theory for neural codes through multiple stages of the visual system and in other sensory systems (Olshausen & Field, 2004) suggests that our theoretical results may have broad application to the role for space-time inseparable sensory stimuli and STDP in the development of their respective circuits.

We have also demonstrated that the mechanism of periodic pattern formation is, in theory, applicable over a broad range of wave speeds and STDP time constants, and that the spatial frequency of the pattern varies in a predictable way with these parameters. This has allowed us to map particular examples of travelling waves onto our theoretical results, which we examine in greater detail in the following chapters of this thesis. In Chapter 4, we describe how periodic patterning may be applied to the development of receptive fields, and explore this further in the context of retinal waves in Chapter 5. In Chapter 6, we take advantage of the dependence of the dominant spatial frequency, k^* , on v and τ_+ to apply the periodic nature of the patterns to the development of simple cell receptive fields in primary visual cortex.

Appendix

Here, we show how firing rate of an input neuron when recruited by a wave, and the excitatory postsynaptic potentials (EPSPs) elicited in the output neuron, enter the expression for the synaptic strength dynamics in Eq. 3.6. These terms do not compromise the assumptions required to derive Eq. 3.3, so we pick up the derivation at this point. The input firing rate and EPSPs enter the derivation in the calculation of the correlation function:

$$C(x, \Delta x) = \int_{-\infty}^{\infty} dt S_{\text{in}}(x, t + \Delta t) S_{\text{out}}(t), \quad (3.11)$$

where $S_{\text{in}}(x, t)$ and $S_{\text{out}}(t)$ are the input and output firing rates, respectively. The wavefront in the input layer is still described by a travelling, $\delta(x - vt)$, as in the main text, but the firing rate of a single input neuron that is recruited by the wavefront is now described by $\alpha(t)$, which is zero valued for $t < 0$. The input firing rate during a wave is the convolution in time of the travelling pulse with $\alpha(t)$:

$$\begin{aligned} S_{\text{in}}(x, t) &= \int_{-\infty}^{\infty} dt' \delta(x - vt') \alpha(t - t') \\ &= \alpha(t - x/v), \end{aligned} \quad (3.12)$$

and the output firing rate is the convolution in time of the input firing rate with the EPSP, $\epsilon(t)$:

$$\begin{aligned} S_{\text{out}}(t) &= \int_{-\infty}^{\infty} dx w(x) \int_{-\infty}^{\infty} dt' \int_{-\infty}^{\infty} dt'' \delta(x - vt'') \alpha(t' - t'') \epsilon(t - t' - t'') \\ &= \int_{-\infty}^{\infty} dx w(x) \int_{-\infty}^{\infty} dt' \alpha(t' - x/v) \epsilon(t - t' - x/v) \\ &= \int_{-\infty}^{\infty} dx w(x) [\alpha(t - x/v) * \epsilon(t - x/v)] \\ &= w(vt) * \alpha(t) * \epsilon(t). \end{aligned} \quad (3.13)$$

where $*$ denotes convolution. Substituting Eqs. 3.12 and 3.13 into Eq. 3.11, the correlation function can be found thus:

$$\begin{aligned} C(x, \Delta t) &= \int_{-\infty}^{\infty} dt \alpha(t + \Delta t - x/v) [w(vt) * \alpha(t) * \epsilon(t)] \\ &= \alpha(\Delta t - x/v) * \alpha(x/v - \Delta t) * \epsilon(x/v - \Delta t) * w(x - v\Delta t). \end{aligned} \quad (3.14)$$

The synaptic strength dynamics can now be computed by inserting Eq. 3.14 into Eq. 3.3:

$$\begin{aligned} \partial_T w(x) &= \eta \int_{-\infty}^{\infty} d\Delta t K(\Delta t) C(x, \Delta t, T) \\ &= \eta \int_{-\infty}^{\infty} d\Delta t K(\Delta t) [\alpha(\Delta t - x/v) * \alpha(x/v - \Delta t) * \epsilon(x/v - \Delta t) * w(x - v\Delta t)] \\ &= \frac{\eta}{v} \int_{-\infty}^{\infty} d\Delta x K(\Delta x/v) [\alpha((\Delta x - x)/v) * \alpha((x - \Delta x)/v) * \epsilon((x - \Delta x)/v) * w(x - \Delta x)] \\ &= \eta K_v(x) * \alpha(-x/v) * \alpha(x/v) * \epsilon(x/v) * w(x), \end{aligned} \quad (3.15)$$

where $K_v(x)$ has been introduced as a rescaled copy of K : $K_v(x) = v^{-1}K(x/v)$, as we had done in the main text. Here, we have arrived at the result expressed in Eq. 3.5. We go on to find solutions for this expression of the synaptic strength dynamics in the main text.

4. Results

Pattern formation as a mechanism to refine receptive fields

Background

The mature state of a neuron's receptive field (RF), and thus its functional role in processing information, depends on several stages of refinement during early development. The early stages of development throughout the brain are largely independent of correlated neural activity. Rather, a scaffold for the initial wiring of neural circuits is established by molecular guidance cues that are expressed in concentration gradients in different brain areas (Kolodkin & Tessier-Lavigne, 2011). By this process, molecular guidance establishes topographic maps, in which nearby neurons in one brain area target nearby neurons in another, connected area. Once a coarse topography is laid out, further refinement takes place by a process that does depend upon correlated activity (Katz & Shatz, 1996). Because of the prevalence of travelling wave activity patterns during these early stages of wiring and topographic refinement, the spatiotemporal correlations induced by travelling waves have been considered to play a role in refinement by a Hebbian process.

A neuron's RF is typically determined by measuring its response to a particular stimulus. Thus, the RF can depend quite heavily on the particular response properties of this neuron's feedforward inputs from the sensory organ, as well as on lateral inputs from other downstream neurons. The simplicity of the model we use throughout this Thesis, in which a layer of input neurons connects via feedforward synapses to a single output neuron, allows us to simplify our definition of the RF as the region of the input layer to which the output neuron is most responsive. Therefore, in our model, the output neuron's RF depends only on the feedforward synaptic strengths.

The dependence of refinement on correlated activity is particularly apparent in the visual system. The most prominent form of topography in the visual system is the retinotopic map, for which adjacent points in visual space, and thus in the retina,

are encoded by adjacent neurons in downstream visual areas. Before the retina is capable of transmitting visual signals downstream, spontaneous waves of spiking activity propagate across the retinal ganglion cell (RGC) layer, driving similar activity patterns downstream (Ackman et al., 2012). Retinal waves progress through three stages, during which waves have distinct properties. Stage 2 waves, which in mammals travel the furthest and at the slowest speed, occur when RGC axons begin to establish and refine arbours in the superior colliculus (SC) (Huberman et al., 2008b). Disrupting the normal propagation of stage 2 retinal waves has a profound effect on the retinotopic map in the SC, resulting in the diffuse and imprecise targeting of RGC axons, enlarged axonal arbours, and RFs in SC that are approximately twice as wide as those in wild type (Simon & O’Leary, 1992a; McLaughlin et al., 2003a; Chandrasekaran et al., 2005; Pfeiffenberger et al., 2006; Huberman et al., 2008a; Dhande et al., 2011).

Using the definition of a RF above, retinotopic refinement can be broken down into two components: 1) the preservation of spatial relationships downstream; 2) the precision with which a single downstream neuron encodes an area of visual space, i.e. its RF. After eye opening, when the retina is fully functional and able to transmit signals from visual stimuli, visual experience becomes important for the maintenance of RFs in both the SC and lateral geniculate nucleus (LGN) (Akerman et al., 2002; Carrasco et al., 2005; Hooks & Chen, 2006; Hooks & Chen). Furthermore, RFs in the primary visual cortex continue to develop and become tuned to a number of visual features, including the spatial frequency, orientation and motion of visual stimuli, in a manner that depends on the spatial and temporal structure of correlations in the visual environment (Hubel & Wiesel, 1963; Braastad & Heggelund, 1985; Chapman & Stryker, 1993; White et al., 2001; Li et al., 2006 ; Li et al., 2008).

In this chapter, we focus on the role played by correlated activity patterns in the

refinement of RFs, where we use the term ‘refinement’ to mean an adaptive change in the connectivity of a feedforward network that results from correlations in the inputs. Our approach is therefore applicable to both topographic refinement and also the functional refinement of single, postsynaptic neurons. A role for correlated activity patterns in RF refinement has been addressed previously in modelling studies, which have utilised the fact that activity patterns, particularly those of travelling waves, induce local correlations in space (Wong et al., 1993). These studies demonstrate how strong correlations between nearby input neurons, and weak correlations between more distant inputs, act through Hebbian plasticity rules and synaptic competition to facilitate strengthening of nearby inputs and suppress strengthening of more distant inputs to a common output neuron (Goodhill, 1993; Miller & MacKay, 1994; Song & Abbott, 2001; Xu et al., 2011). Spike-timing dependent plasticity (STDP) is one such form of Hebbian plasticity, but which is sensitive to both spatial and temporal correlations (Kempster et al., 1999; Song & Abbott, 2001). In Chapter 3, we demonstrated how STDP interacts with travelling wave activity to build spatially periodic connectivity patterns into a feedforward network. It is therefore worth considering how both the spatial and temporal correlations induced by travelling waves may contribute to RF refinement.

Here, we build upon our findings in Chapter 3 to address how travelling waves may refine or reshape a RF through STDP, by virtue of their space-time inseparable correlations. In particular, we consider how the spatial frequency of the periodic connectivity pattern described in Chapter 3 may be imposed on an immature RF to modify its size and shape. We extend our findings to waves that travel in two dimensions, and demonstrate that our theoretical findings in Chapter 3 are upheld even with complex, noisy wave patterns.

Methods

Travelling plane waves in the input layer

Unless otherwise stated, waves were made to travel with speeds of 4 mm/s. Neurons that were recruited by a wave always had firing rates of 50 Hz. Otherwise, when an input neuron was not being recruited by a wave, it was silent. In 1D simulations, consecutive waves were always made to travel in alternating directions so as to prevent drift of the connectivity patterns that developed. In 2D simulations, waves travelled in one of the following set of directions: alternating 0° and 180° directions; alternating 90° and 270° directions; one of sixteen, equally spaced directions around the compass, following a random order.

Complex wave patterns

To generate complex and noisy waves in 2D, we implemented a simulation based on a previously published model of the developing mammalian retina (Feller et al., 1997). A full description of the retina model is provided in General Methods. Here, we describe specific modifications made for the simulations in this chapter.

The complex wave patterns generated by the Feller et al. model travel at speeds between 0.1–0.3 mm/s, which reflect experimentally observed retinal waves. In order to produce complex wave patterns that propagated with speeds that are comparable to those used in the previous chapter and for the plane waves in this chapter, which are an order of magnitude faster, we scaled time in the retina model by a constant factor, F_t . However, we showed in Chapter 3 that the burst duration could significantly affect the outcome of patterning in the network’s connectivity, and therefore required that it remained unchanged by the rescaling of time. To achieve this, we isolated every burst of spikes generated by the retinal wave model for each neuron

and multiplied the burst onset time by $F_t = 0.21$ for slow waves, 0.11 for medium waves and 0.07 for fast waves. Their respective wave speeds are stated in the results. From every new burst onset time, we used a Bernoulli process to generate a burst of spikes lasting 100 ms with a firing rate of 50 Hz at a 1 ms temporal resolution. If a burst temporally overlapped with the following burst, the following burst was deleted. In this way, the wave trajectory was very well maintained by rescaling time, with no affect on the burst duration or mean firing rate within a burst.

We ran three repeated trials of these wave simulations using different seeds for the random number generator, as described in General Methods, such that each trial produced a different sequence of wave patterns.

Receptive fields in the spiking neuron simulation

In all simulations in this chapter, we set the synaptic strengths, $w_0(x)$, such that the output neuron receives strong synaptic input from only a small region of the input layer, which we call the output neuron’s receptive field (RF). A schematic of the network architecture is provided in Fig. 4.1A. RFs were constructed such that, for a RF size of RF_0 , synapses within a distance of $RF_0/2$ from the centre of the input layer had maximum strength and all other synapses had zero strength. The spatial profile of initial synaptic strengths is drawn in grey in Fig. 4.1B. This is a very different scenario from Chapter 3, in which the initial synaptic strengths were randomly distributed about a constant value. Under that condition, $w(x)$ would become periodic in space with a spatial frequency of k^* as $T \rightarrow \infty$. Here, in simulations where $w_0(x)$ had the RF structure, we found that additional peaks in the connectivity would eventually emerge until a periodic pattern occupied the entire network, which falls short of the notion of a RF. We therefore took advantage of the fact that, in biological neural circuits, restrictions on the size of a neuron’s RF are imposed by the spatial extent of axonal arbores providing input and the

size of its own dendritic arbour receiving that input. In the simulations of this chapter, we imposed an arbour function, $A(x)$, which prevented synapses beyond the arbour from ever developing a connection with the output. The arbour function was implemented by multiplying $A(x)$ with $\partial_T w(x)$, where

$$\begin{aligned} A(x) &= 1, \text{ if } |x - x_0| \leq a_0 \\ &= 0, \text{ if } |x - x_0| > a_0 \end{aligned}$$

with x_0 the position of the RF centre and a_0 the radius of the arbour.

Results

Travelling waves are instructive for shaping receptive fields via spike-timing dependent plasticity

In this chapter, we examine whether the interaction of travelling waves and STDP rules, which can impose spatially periodic connectivity patterns on a uniform field of synaptic strengths (see Chapter 3), might be a plausible mechanism for shaping RFs. The RFs we model (see Methods), which restrict synaptic strengths to the range $[0, 1]$, do not lend themselves easily to analytical investigation. Therefore, we began by using the simulations of spiking neurons as in Chapter 3, but with different initial distributions of synaptic strengths. During development, coarse RFs are thought to be set up by molecular cues before activity dependent refinement takes place (Huberman et al., 2008b). Therefore, at the beginning of the simulation, we set synapses connecting the centre of the input layer to maximum strength and the remaining synapses to zero, thus endowing the output neuron with an initial RF. The diameter of the initial RF was set by the parameter RF_0 .

In a first set of simulations, we set $RF_0 = 0.80$ mm, which corresponds, for example, to the approximate retinal projection of a stimulus that matches the RF size of immature simple cells in 4 week old cats (Braastad & Heggelund, 1985; DeAngelis et al., 1993a), immature neurons in lateral geniculate nucleus of ferrets at eye opening (Tavazoie & Reid, 2000; Akerman et al., 2004), and poorly refined RF sizes in SC of adult mice that have had disrupted stage 2 retinal waves (Chandrasekaran et al., 2005). To calculate RF_0 , we converted measured RF sizes in the experiments cited using the conversion formula in General Methods.

By interpreting the RF as a single cycle in the type of periodic patterns obtained in Chapter 3, where RF_0 is some fraction of the pattern’s wavelength, we were able to examine how properties such as the wave speed and the STDP time constant influence RF development, due to the effect these variables have on the dominant spatial frequency of the connectivity pattern, as shown in Chapter 3. Therefore, we ran a set of simulations between which the STDP time constant and wave speed were varied. In Figs. 4.2A–C (centre panels), we show that, with wave speeds of 4 mm/s and a STDP rule with $\tau_+ = 20$ ms, the RF can be refined to a smaller size. This is because repeatedly convolving the RF with this particular $\kappa(x)$ strengthens the higher spatial frequencies and suppresses the lower spatial frequencies of the RF, effectively chipping away at its edges. We examined how the dominant spatial frequency of $\kappa(x)$ determined RF development by changing the time scales of STDP or by changing the wave speed. In Fig. 4.2A, we show the development of RFs for different STDP time scales. We increased the dominant spatial frequency by reducing the time scale for synaptic weakening such that $\tau_- = \tau_+ = 20$ ms, as in Song et al. (2000), and setting $A_- = 1.1$ so as to maintain the constant negative value of the integral, $\int_{-\infty}^{\infty} d\Delta t K(\Delta t)$, and thus maintain competition between synapses. Under this condition, the RF split into subfields (Fig. 4.2A, left panel), which were consistent with multiple cycles of the periodic pattern obtained from simulations

in Chapter 3. A low spatial frequency was achieved by setting $\tau_+ = 50$ ms and $\tau_- = 100$ ms ($A_- = 0.55$), causing the RF to grow in size (Fig. 4.2A, right panel). RFs were split into subfields or expanded in a similar fashion by decreasing the wave speed to $v = 2$ mm/s (Fig. 4.2B, left) or by increasing it to $v = 10$ mm/s (Fig. 4.2B, right), respectively. These results suggest that the characteristic spatial frequency of $\kappa(x)$, or rather the corresponding wavelength, relative to the size of the initial RF, determines how the RF is modified. To confirm this, we varied RF_0 , keeping $\kappa(x)$ fixed ($\tau_+ = 20$ ms, $\tau_- = 40$ ms, $v = 4$ mm/s). When RF_0 was narrow ($RF_0 = 0.28$ mm), the RF expanded (Fig. 4.2C, left), and when RF_0 was wide ($RF_0 = 1.4$ mm), the RF split into subfields (Fig. 4.2C, right). Thus, the way in which the RF is modified depends on the shape of $\kappa(x)$ relative to the initial receptive field.

In Chapter 3, $\kappa(x)$ set the spatial frequency of the periodic connectivity pattern. By interpreting a RF as a single cycle of a periodic pattern, we expected $\kappa(x)$ to also set the final size of RFs (or subfields) in our simulations. To test this, we varied RF_0 over a wide range between simulations. In Fig. 4.2D, we plot the number of strong synapses (with strength > 0.5) over the duration of simulations for the RF_0 values tested. RFs either maintained a single subfield, or split into two or three subfields as RF_0 was increased (purple, yellow and blue traces in Fig. 4.2D, respectively). Regardless of RF_0 , the mean final size of each subfield was nearly identical. In other words, two or three times as many synapses stabilised with maximal strength when the RF split into two or three subfields, respectively, compared with when the RF maintained just a single field. Increments of the mean number of synapses per subfield (for RFs with three subfields) are drawn on the right of Fig. 4.2D to illustrate this result. Small biases were evident, however, whereby larger values of RF_0 tended to give rise to larger RFs or subfields. This considered, the final number of strong synapses were easily distinguishable between cases in which one, two or three




Figure 4.1: **Network architecture and dynamics.** **A.** The model used in this chapter is a simple modification to that of Chapter 3. In this case, simulations begin with an architecture that endows the output neuron with a RF, whereby synapses within a radius of RF_0 from the centre have maximal strength (solid grey connectors), and synapses beyond this radius have zero strength. An arbour function, with a radius that exceeds RF_0 , prevents the development of synapses outside the arbour, which are effectively disconnected from the output. Synapses that are initialised with zero strength but lie within the arbour can be modified during the simulation (dashed grey connectors). Thus, the output neuron only receives synaptic input when the wave recruits input neurons from within the arbour. **B.** Grey: initial synaptic strengths, $w_0(x)$, as a function of the input location. Green: arbour function. Both plots are scaled in space so as to align with the corresponding components of the diagram in **A**.




Figure 4.2: Refinement of receptive fields. **A-C.** Examples of RF development over time as different parameters in the simulation are varied. Green lines mark the boundary of the arbour function. **A.** Changes to STDP time scales. Left to right: $\tau_+ = 20$ ms, 20 ms and 50 ms, $\tau_- = 20$ ms, 40 ms and 100 ms, $A_- = 1.1$, 0.55 and 0.55. Fixed values: $RF_0 = 0.8$ mm, $v = 4$ mm/s. **B.** Changes in wave speed. Left to right: $v = 2$ mm/s, 4 mm/s and 10 mm/s. Fixed values: $RF_0 = 0.8$ mm, $\tau_+ = 20$ ms, $\tau_- = 40$ ms, $A_- = 0.55$. **C.** Changes in the initial RF size. Left to right: $RF_0 = 0.28$ mm, 0.8 mm and 1.4 mm. Fixed values: $v = 4$ mm/s, $\tau_+ = 20$ ms, $\tau_- = 40$ ms, $A_- = 0.55$. A large arbour had to be used when $RF_0 = 1.4$ mm in order to accommodate this larger initial RF size. **D.** Development of the mean number of synapses with strength > 0.5 in simulations using RF_0 values from 0.2 mm to 3.2 mm in steps of 0.2 mm. Traces tend towards one of three values, depending on the number of subfields that develop. Blue: RFs split into three subfields; yellow: RFs split into two subfields; purple: RFs maintain just a single field.

subfields developed. Thus, the size and shape of the final RF tightly corresponds with the wave, STDP and initial RF properties. In the following results sections, we explore this relationship in greater detail and, in the Discussion, suggest a role for it in several stages of visual development.

A quantitative relationship between periodic patterning and refinement of receptive fields

The properties of both travelling waves and mechanisms of plasticity vary considerably between animals, between synapses and over time. In Chapter 3, we demonstrated the robustness with which periodic connectivity patterns would develop in the presence of several types of noise and variability that is common to travelling wave activity. Here, along similar lines, we examine the robustness of RF refinement in terms of the range over which model parameters satisfy conditions for refinement. In doing so, we draw more concrete parallels between the construction of periodic connectivity patterns and RF development. To do this, we computed numerical solutions to Eq. 3.6 in Chapter 3 for pairs of v and τ_+ values over the wide range of wave speeds and STDP time constants used to generate Fig. 3.3 in Chapter 3. For each v and τ_+ pair, we set the initial synaptic strengths, $w_0(x)$, to be identical to that in the above simulations with a fixed $RF_0 = 0.8$ mm. To further align numerical solutions with the above simulated results, we applied nonlinearities such as bounding of synaptic strengths in the range $[0, 1]$, changes in wave direction and an arbour function during each iteration (see General Methods). In Fig. 4.3A, we plot the evolution of $w(x)$ for a RF that was computed numerically with parameters chosen to match those of the central panels in Figs. 4.2A–C. Using these solutions for $w(x)$, we constructed a phase diagram for the size and shape of RFs that developed, which we call the refinement phase space. For each v and τ_+ pair, we counted the number of subfields in the final RF and, for solutions that did not divide into multiple subfields, measured the change in RF size as a percentage of RF_0 . In other

words, we hereafter focus our analysis on the change in size of RFs consisting of a single contiguous patch of inputs. In particular, we highlight those cases in which RFs are refined to smaller sizes, which is a common phenomenon in developing RFs of the visual system. In Chapter 6, we return our attention to the division of RFs into periodic subfields in relation to the development of simple cells in primary visual cortex.

In Fig. 4.3B, we plot the refinement phase space. The area coloured in blue corresponds to final RFs that consisted of a single field and had decreased in size relative to $RF_0 = 0.8$ mm. Darker shades of blue indicate a greater decrease in size. Overlaid are the iso-frequency contours of the predicted spatial frequencies, k^* , as in Fig. 3.3 in Chapter 3. The boundaries of the phase space, in which RF sizes decreased, closely follow the iso-frequency contours, with greater decreases (darker shades) corresponding to higher spatial frequencies. In the region labelled α (unshaded), at which $\kappa(x)$ has even higher characteristic spatial frequencies, RFs split into subfields. In the region labelled β (unshaded), at which spatial frequencies are lower, RFs expanded. Because the outcome of the final RF depended on the initial RF size, we repeated numerical solutions for the refinement phase space, setting $RF_0 = 0.44$ mm, which just exceeded the size of the smallest contiguous RF in solutions when $RF_0 = 0.8$ mm. For solutions when $RF_0 = 0.44$ mm, the region of phase space in which RFs maintained a single field and decreased in size (coloured purple) corresponded to lower wave speeds and shorter STDP time constants, i.e. higher spatial frequencies in $\kappa(x)$. Taken together, these results suggest that smaller initial RF sizes require higher spatial frequencies in $\kappa(x)$ to achieve a particular pattern of RF refinement. The area of the blue and purple regions of phase space also show that, in order for RFs to be refined to a smaller size by the asymmetric STDP rule, the maximum possible wave speed can be just over a factor of two times greater than the minimum possible wave speed. At least the same range in τ_+ can also

enable such refinement of RFs. Thus, the mechanism of RF refinement described in this chapter allows for some variation in STDP properties between synapses, variation in wave propagation, or even fluctuations in both of these phenomena over time.

An interesting corollary of these results is that changes in wave speed and STDP properties can follow changes in RF size, thus allowing continual refinement of a RF beyond that of a fixed v and τ_+ combination. To exemplify this point, we ran a simulation using the same parameters as those used for the central panels in Figs. 4.2A–C. However, once the RF had refined to a smaller size, we decreased the wave speed from 4 mm/s to 2.5 mm/s. Note that, according to Fig. 4.3B, $v = 2.5$ mm/s and $\tau_+ = 20$ ms would cause a RF with $RF_0 = 0.8$ mm to split into subfields, as this point in the phase space lies to the lower left of the blue shaded area. However, as shown in Fig. 4.3C, an initial period of refinement, which maintained a single field with $v = 4$ mm/s, set up further refinement via a decrease in wave speed to $v = 2.5$ mm/s without splitting the RF into subfields. Increasing the wave speed back to 4 mm/s at a later time caused the RF to expand and return to its previous size. By the same reasoning, changes in the time constants of STDP should have the same effect.

These results correspond exactly with our results from the simulations in Fig. 4.2, confirming our hypothesis that $\kappa(x)$ and RF_0 jointly determine the final shape and size of the RF. This was not unique to the asymmetric STDP rule. In Fig. 4.3D, we show a similar correspondence between refinement and the iso-frequency contours for a symmetric STDP rule. However, we found that, in order for a RF to be refined to smaller sizes without splitting into subfields, the STDP rule must exhibit a bias towards synaptic weakening. In Fig. 4.3E, we plot the refinement phase space of a 0.8 mm RF for three asymmetric STDP rules that differed only in A_- , which scaled the amplitude of the negative lobe in the rule. When $A_- = 0.5$, we have




Figure 4.3: **Relationship between periodic patterning and receptive field refinement.** **A.** Example of RF refinement obtained by numerically solving Eq. 3.6 in Chapter 3. Green lines delimit the arbour function. **B.** Refinement phase space for RFs as a function of wave speed and STDP time constants. Coloured regions: areas of phase space in which RFs maintain a single field that is refined to smaller sizes; blue: $RF_0 = 0.8$ mm; purple: $RF_0 = 0.44$ mm. Shading indicates the total change in RF size as a percentage of RF_0 . Region α corresponds to higher spatial frequencies in which RFs split into subfields. Region β corresponds to lower spatial frequencies in which RFs expand. Overlaid are the iso-frequency contours of Fig. 3.3 from Chapter 3. Solid contours correspond to spatial frequencies with integer exponents of 10 (lower left to upper right: 10 cyc./mm, 1 cyc./mm and 0.1 cyc./mm). **C.** Incremental changes in wave speed can continually refine RFs while maintaining a single subfield structure. Shown is the development of a RF, averaged over 16 trials, during which the wave speed is first set to 4 mm/s and then decreases to 2.5 mm/s, after which the RF decreases in size. Increasing the wave speed back to 4 mm/s returns the RF to its previous refined size. Purple trace: plot of the mean number of synapses with strength > 0.5 as an indicator of RF size. **D.** Refinement phase space using a symmetric STDP rule. RF_0 is the same here as it is for the blue region in panel **B**. Numbers next to the iso-frequency contours indicate their spatial frequency in cycles/mm. **E.** Refinement phase spaces, using the asymmetric STDP rule, for different biases towards synaptic weakening. Left: $A_- = 0.50$, which corresponds to $\int_{-\infty}^{\infty} d\Delta t K(\Delta t) = 0$. Centre and right: $A_- = 0.55$ and 0.60 , respectively, which corresponds to $\int_{-\infty}^{\infty} d\Delta t K(\Delta t) > 0$. Red circles lie at the same v and τ_+ coordinates, to compare the effect of different A_- values on RF refinement. The refinement phase space moves to lower spatial frequencies as A_- increases.

$\int_{-\infty}^{\infty} d\Delta t K(\Delta t) = 0$, such that synapses no longer compete for connections with the output. Consequently, increasing the spatial frequency of $\kappa(x)$ would not decrease the size of the RF before causing it to divide into subfields (left panel). However, increasing A_- , such that $\int_{-\infty}^{\infty} d\Delta t K(\Delta t) < 0$, did allow for refinement of a RF to smaller sizes. Furthermore, increasing A_- ($A_- = 0.55$, middle panel; $A_- = 0.6$, right panel) moved the refinement phase space, making refinement possible at lower characteristic spatial frequencies and, therefore, larger v and τ_+ values. As such, for a given v and τ_+ pair (red circles in Fig. 4.3E), greater refinement was achieved with increased bias for synaptic weakening. Equipped with a basic understanding of the relationship between periodic patterning and RF refinement, we examine in the next section how travelling waves and STDP impact the development of RFs in two dimensions.

Refinement of receptive fields in two spatial dimensions

In the developing brain, travelling waves frequently propagate via horizontal connections that extend parallel to the cortical surface, cerebellar surface or retina, and their trajectory should therefore be considered as two dimensional (2D). In two dimensions, the trajectory of wavefronts can exhibit considerably greater complexity than in one dimension. In order to later tackle the influence of realistic and complex waves on RF development, we first demonstrate differences between the 1D and 2D cases. The variables x and v in Eq. 3.6 of Chapter 3 can now be considered as 2D vectors, $\mathbf{x}$ and $\mathbf{v}$, in the x - y plane. As such, we expect the connectivity pattern to adopt the characteristic spatial frequency of $\kappa(\mathbf{x})$, $\mathbf{k}^*$, and thus for RFs to undergo refinement, but only along a direction parallel to $\mathbf{v}$. To test this, we generated wave stimuli in a 2D input layer consisting of 64×64 units arranged on a square lattice. Plane waves were generated in a similar fashion to the 1D scenario, moving with a constant speed in alternating directions, from one side of the input layer to the opposite side with a wave speed of 4 mm/s and an asymmetric STDP rule using

$\tau_+ = 20$ ms. The single output neuron was provided with an initial, circular RF 0.8 mm in diameter (Fig. 4.4A, left panel). This choice of parameters, as shown in Fig. 4.3B, should refine the RF to a smaller size. In Fig. 4.4, we show the final, stabilised RFs. When plane waves in the input layer travelled along a horizontal direction, the RF was refined along the horizontal axis, as expected (Fig. 4.4A, top panels). However, along the vertical axis, the RF grew in size until it spanned the full extent of the arbour from top to bottom. This is in line with the fact that the wave speed, v , now depends on direction. Thus, parallel to the direction of wave propagation, v is simply the wave speed quoted, 4 mm/s. In the orthogonal direction, v is infinite, resulting in a periodic pattern with a spatial frequency of zero. Thus, when waves instead travel along the vertical direction, the RF refines along the vertical axis and grows along the horizontal axis (Fig. 4.4A, bottom panels).

Following from these results, it was expected that RFs would maintain their circular symmetry under the influence of waves travelling in multiple directions, but it was not obvious how or whether RFs would refine, as RF narrowing along one axis is concomitant with expansion along the orthogonal axis. To address this, using the same parameters as above, we ran simulations in which each wave was randomly assigned one of sixteen possible directions, equally spaced around the compass. The final RF for this case is drawn in Fig. 4.4B (left panel). Under these conditions, RFs maintained an approximate circular symmetry but the RF area increased in size (Fig. 4.4C, dark blue). Nevertheless, refinement to smaller RF sizes could be achieved by changing appropriate parameters, such as the wave speed or the ratio of synaptic weakening to strengthening, as explained in the results of Fig. 4.3. Thus, increasing the dominant spatial frequency, k^* , by decreasing the wave speed to 3 mm/s (Fig. 4.4B middle panel), or increasing the amplitude for synaptic weakening in the STDP rule from 0.55 to 0.6 (Fig. 4.4B right panel), both gave rise to final RFs that were smaller than the initial RF (Fig. 4.4C, light blue and purple,




Figure 4.4: Refinement of receptive fields with 2D plane waves. **A**. Refinement depends on the direction of wave propagation. Examples of refinement for two RFs are given here (top row and bottom row, respectively; increasing time from left to right). If an initial RF is circular (left), then waves travelling along a horizontal direction will refine the RF along the horizontal axis only (top row). The RF necessarily expands to the arbour boundary along the orthogonal axis. Likewise, waves travelling along a vertical direction refine RFs along the vertical axis only (top row). Increasing time from left to right. Arrows: direction of wave propagation. Solid green circle: arbour boundary. Dashed green circle: boundary of the initial RF. **B**. Final RFs that were exposed to plane waves travelling randomly in 16 possible directions. Left: using the same wave speed (4 mm/s) and STDP parameters ($\tau_+ = 20 \text{ ms}$, $A_- = 0.55$) as in **A**. Centre: using a slower wave speed (3 mm/s) to lower the characteristic spatial frequency of $\kappa(\mathbf{x})$. Right: using the faster wave speed (4 mm/s) but shifting refinement phase space to lower frequencies by increasing the bias for synaptic weakening ($A_- = 0.60$). **C**. Summary of the development of RF sizes for each of the conditions in **B**, averaged over three trials for each condition.




Figure 4.5: Properties of complex, noisy waves generated using the modified Feller et al. model. **A.** Snapshots of two example waves generated by this model, with time increasing from left to right in steps of 0.4 s. Dashed orange curves mark the spatial boundaries within which activity was assigned to the same, isolated wave. Black pixels mark input neurons that fired at least one spike during a 0.1 s interval. Grey pixels mark input neurons that fired at least one spike outside of the wave boundary. The wave isolation algorithm excluded activity outside the boundary from further analysis. **B** sCCGs for the two example waves in **A**. Darker colours indicate stronger correlation, which approximately follow a ridge that extends from the origin. Purple lines denote the linear regression that was fit to the crest of the ridge. The inverse of the gradient of the linear regression provided an estimate of the speed at which correlations were propagated across the network. The speeds calculated for these two waves are given at the top of each sCCG. **C.** Distributions of wave speeds for waves taken from the same original simulation that was rescaled in time by constant factors, F_t . Each distribution corresponds to a different value of F_t . Waves were isolated for analysis over a 2000 s period for a given F_t . Median wave speeds are indicated by arrows above each distribution. The number of analysed waves populating each distribution are given next to the arrows.

respectively). This shows that the principle of using STDP and travelling waves to refine a RF extends from 1D to 2D networks for simple, idealised waves. In the following section, we will examine how this holds up when waves are more irregular and realistic.

Refinement of receptive fields by complex wave patterns

In the early visual system, retinal wave activity is required for the early refinement and pruning of retinofugal efferents to the superior colliculus (McLaughlin et al., 2003a; Chandrasekaran et al., 2005; Pfeifferberger et al., 2006; Huberman et al., 2008a; Dhande et al., 2011). At a later stage, visual experience is necessary for the maintenance of RFs in SC (Carrasco et al., 2005) and of the total synaptic input to neurons in the lateral geniculate nucleus (LGN) (Hooks & Chen, 2006; Hooks & Chen), as well as the normal development of mature RF properties in primary visual cortex, which involves a refinement in size and spatial frequency tuning (Braastad & Heggelund, 1985; DeAngelis et al., 1993a; Fagiolini et al., 1994). Both retinal waves and natural scenes contain very complex features, for which the underlying spatiotemporal statistics have been hypothesised to play a role during different stages of visual development. Common to activity patterns resulting from both retinal waves and natural scenes is motion. However, the precise statistics of retinal waves and natural scenes are difficult to quantify and are likely to differ between animals. Despite this, several studies have provided insights to particular cases (Feller et al., 1997; Ford et al., 2012; Dong & Atick, 1995; Reinagel & Zador, 1999). We therefore ran simulations to test whether complex patterns of wave activity could give rise to refinement of RFs in the same manner as we have thus far demonstrated using simple plane waves.

Here, we have used a modified version of the retinal wave model from Feller et al. (1997) (see Methods) to generate complex patterns of activity in the input layer,




Figure 4.6: Refinement of RFs by complex wave patterns. **A**. Example final RFs that developed under different wave speed and τ_+ combinations. Centre: under the influence of medium speed waves, with $\tau_+ = 20 \text{ ms}$, RFs maintained a single field that was refined to a smaller size. Top: slow waves refined RFs to even smaller sizes and, in this example, caused the RF to split into two subfields (cf. Figs. 4.2A). Bottom: fast waves caused RFs to expand beyond their original size. Left: with $\tau_+ = 15 \text{ ms}$, RFs refined to even smaller sizes. Right: with $\tau_+ = 25 \text{ ms}$, RFs refined to a smaller size but less than when $\tau_+ = 20 \text{ ms}$. **B**. Summary of the development of RF areas under the influence of slow, medium and fast waves, averaged over three trials. Dashed grey line: initial RF area. **C**. Summary of the development of RF areas in simulations using different STDP time constants, averaged over three trials. Dashed grey line: initial RF area. RF areas were calculated by multiplying the number of synapses with strength > 0.5 by the input cell spacing ($34 \mu\text{m}$) squared.

leaving the details of the precise statistical nature of retinal waves and natural scenes for future work. The Feller et al. model generates waves via a 2D grid of cells that are connected laterally by excitatory synapses. Stochastic self-excitation causes these cells to become active spontaneously, providing excitatory synaptic input to nearby cells and triggering cascades of excitatory synaptic transmission that propagate throughout the network, while strong refractory currents quickly return cells to the rest state. As such, wavefronts of activity propagate along winding trajectories between refractory domains of the network. These wavefronts were then filtered through a second layer of spiking units that provided the input for simulations of RF development.

Start with saying that waves speeds could not be controlled due to the stochastic process of wave generation. Then say how we had to isolate waves in order to measure them. Then how we measured them using the sCCG. Then how we varied wave speeds in different sets of simulations.

To test whether complex waves could refine RFs by the process of pattern formation shown above, we varied the wave speed and the STDP time scale between simulations. However, due to the stochastic nature of wave generation and propagation in the model, we could not control the speed of individual waves. While parameters of the retinal wave model can be modified to change the average wave speed (Feller et al., 1997; Butts et al., 1999), such changes have the undesirable consequence of altering the spatiotemporal structure of waves and the pattern of wave trajectories, which would confound any influence the wave speed had over the development of RFs. Therefore, to generate wave patterns with different overall wave speeds, while maintaining their spatiotemporal structure, we simulated long periods of wave activity and rescaled these activity patterns in time by a factor, F_t . At the same time, we ensured that every burst in each unit had a duration of 0.1 s and mean firing rate of 50 Hz (full details are in Methods). In this way, average wave speeds could be

scaled by $1/F_t$. In order to measure the wave speed, individual wave events had to be isolated. Briefly, this involved tracking the centre of mass (COM) of firing rates for each contiguous area of active units, assigning domains with COMs that were nearby in space and time to the same wave (see General Methods). Two examples of isolated waves that were generated by the modified Feller et al. model are provided in Fig. 4.5A. Drawn in orange are the boundaries within which activity was assigned to the same wave. We measured the speed of isolated waves by computing the relationship between their spatial and temporal correlations in a space-time cross correlogram (sCCG, General Methods). The sCCG maps the mean correlation strength between the spiking activity of all input neuron pairs as a function of their spatial separation, Δr , and the temporal lag between their spikes, Δt . sCCGs for the two example waves are provided in Fig. 4.5B. For travelling waves, the sCCG produces a ridge of strong correlations, whereby the Δr and Δt coordinates along the crest of the ridge have an approximately linear relationship. The wave speed was therefore computed by fitting a linear regression to these coordinates (purple lines in Fig. 4.5B) and taking the inverse of the gradient of the fit. We used three different values of F_t (see Methods) to produce simulated complex wave patterns with either slow, medium or fast wave speeds. In Fig. 4.5C, we plot the wave speed distributions, which have median wave speeds of 0.64 mm/s (slow), 1.25 mm/s (medium) and 1.79 mm/s (fast), corresponding to the three different F_t values.

Using these rescaled wave patterns as input to simulations of RF development, we found that, quite remarkably, the complex wave patterns could shape and refine RFs in much the same manner as simple plane waves. Fig. 4.6A provides examples of final RFs for different v and τ_+ combinations. Using an asymmetric STDP rule with $\tau_+ = 20$ ms, medium speed waves were able to refine the RF to smaller areas (Fig. 4.6A, middle panel). With slow waves, RFs refined to even smaller areas, and in some cases, split into subfields (Fig. 4.6A, top panel). The fast waves, however,

caused RFs to expand over a greater area (Fig. 4.6A, bottom panel). A summary for the evolution of RF areas under these three wave speeds, averaged over repeated trials, is provided in Fig. 4.6B. Similarly, STDP rules with shorter ($\tau_+ = 15$ ms) and longer ($\tau_+ = 25$ ms) time constants produced final RFs with smaller or larger areas, respectively, compared with STDP rules with $\tau_+ = 20$ ms. These effects are summarised in Fig. 4.5E. Thus, these results recapitulated the dependence of RF development on wave speed and STDP time constants, even when the input wave patterns followed complex and noisy trajectories.

Discussion

RFs undergo fine scale refinement in an activity dependent manner during early stages of neural development. This process can be understood for activity patterns that are correlated over small spatial and temporal scales, operating through Hebbian mechanisms of plasticity, whereby strong correlations between nearby inputs facilitate the strengthening of their connections to an output, while competing with and suppressing more distant inputs that are less well correlated. Over short time scales, travelling waves induce strong correlations between inputs over small distances. However, travelling waves can also induce correlations that are equally as strong over large distances and long time scales. Here, we have utilised this fact and demonstrated a novel process by which travelling wave activity patterns may refine RFs through Hebbian plasticity. For STDP rules that specify synaptic weakening over longer time scales than for synaptic strengthening, travelling wave correlations cause nearby synapses to facilitate strengthening and more distant synapses to compete. Thus, RF refinement as we have described in this chapter is not due to the global suppression of weakly correlated inputs, but is due to the interaction of strong wave-induced correlations and STDP. Moreover, this process imposes a characteristic wavelength on the RF (as in Chapter 3). The wave speed and time

scale of STDP are two factors that determine the wavelength, such that changes to these properties have a strong influence over the developmental trajectory of the RF, which depends on the ratio of the wavelength to the initial RF size. We identified three such trajectories: 1) if the wavelength is suitably tuned, the RF can be refined to a smaller size; 2) shorter wavelengths, which result from slower wave speeds or faster STDP decay times, divide the RF into subfields; 3) for longer wavelengths, which result from faster wave speeds or slower STDP decay times, the RF grows to a larger size.

Suitability of spontaneous travelling waves for refinement in different brain areas

Our findings implicate the potential importance of wave and plasticity properties in the early refinement of neural circuits. Different wave speeds are reported in different areas of the brain, such as the cerebellum (Watt et al., 2009), hippocampus (Leinekugel et al., 1998), ventral and dorsal cortex (Conhaim et al., 2010) and the retina (Feller et al., 1997), and are summarised in Table 4.1. If waves in these areas are to play a role in topographic refinement, what RF size would be predicted by our model to develop as a result of these wave speeds, and how does this size relate to the size of the brain area over which the waves propagate? Let us assume that refinement results from an asymmetric STDP rule with a time constant of $\tau_+ = 20$ ms, as is frequently reported (Caporale & Dan, 2008). We can then calculate the predicted spatial frequencies, or rather the corresponding wavelengths, λ , that would emerge in connectivity patterns as a result of these waves, as we had done in Chapter 3 to make Fig. 3.3. The expected wavelengths for waves in these different areas, and the maximum distance over which they can propagate, are also summarised in Table 4.1. In the cerebellum, waves travel a distance of ~ 0.8 - 1.0 mm along the sagittal plane of the folia, and we calculate the wavelength of the connectivity pattern to be $\lambda = 0.8$ mm; in the hippocampus, waves travel a distance of ~ 5 mm along the

septotemporal axis, and $\lambda = 2$ mm; in ventral cortex, waves travel a distance of ~ 1.7 mm around the coronal plane, and $\lambda = 2$ mm; in dorsal cortex, waves travel a distance of ~ 2.9 mm around the coronal plane, and $\lambda = 4$ mm. It is interesting to note that, for travelling waves observed in the cerebellum (Watt et al., 2009) and ventral cortex (Conhaim et al., 2010), and to some extent the dorsal cortex (Conhaim et al., 2010), the characteristic wavelengths of the periodic connectivity pattern that our theoretical results predict are of a similar size to the maximum possible distance travelled by the wave in each respective brain area. The distance travelled by the wave and the characteristic wavelength of the connectivity pattern are the spatial analogues of the duration of the wave and the time scale of STDP, respectively. Therefore, the similarity between the characteristic wavelength and the size of the brain area in which waves can travel challenges the assumption, stated in Chapter 3, that the wave duration must be sufficiently longer than the time scale for STDP. This would suggest that the waves reported in these areas are unlikely to contribute to the refinement of topographic maps through STDP in the manner described in this chapter. However, this does not exclude these waves from giving rise to refinement through alternative Hebbian mechanisms (Goodhill, 1993; Miller & MacKay, 1994).

Activity patterns in the retina might be better suited to our description of RF refinement. In mice, the distance from one edge of the retina to the furthest edge on the other side of the retina is approximately seventy five times larger than the predicted wavelength of the resulting connectivity pattern. Spontaneous retinal waves are prevalent for several weeks prior to the onset of vision and coincide with a period of large scale refinement of retinal afferents in the superior colliculus (SC) (Simon & O’Leary, 1992a; McLaughlin et al., 2003a; Chandrasekaran et al., 2005; Pfeifferberger et al., 2006; Huberman et al., 2008a; Dhande et al., 2011). Imposing a wavelength onto the spatial structure of retinocollicular afferents is one way in

Brain Area	Wave Speed (mm/s)	λ (mm)	Max. propagation distance (mm)
Cerebellum [1]	3	1	1
Dorsal cortex [2]	17	4	3
Ventral cortex [2]	7	2	2
Hippocampus [3]	8	2	5
Retina [4]	0.2	0.04	3

Table 4.1: **Summary of different travelling wave phenomena in the brain.** For each brain area, the typical wave speed is given as well as the expected wavelength, λ , of a connectivity pattern that would emerge due to the interaction of those waves with an asymmetric STDP rule for which $\tau_+ = 20$ ms. Also provided are the maximum distances that waves in these brain areas are able to propagate. Only in the retina is the maximum distance much larger than the expected wavelength. [1] Watt et al., 2009; [2] Conhaim et al., 2010; [3] Leinekugel et al., 1998; [4] Feller et al., 1997.

which retinal waves may drive this refinement. However, retinal waves generate long bursts of spikes, lasting several seconds, which we have shown in Chapter 3 to be ineffective at building periodic connectivity patterns under STDP rules with 20 ms time scales. The application of our analytical results to retinal waves therefore requires consideration of a plasticity rule that is appropriate for long burst durations. Moreover, refinement of RFs by slow retinal waves requires STDP rules to operate over time scales of seconds rather than tens of milliseconds. We address this issue in Chapter 5, and provide direct examples of how the distance propagated by a wave impacts RF refinement.

A role for visual experience in receptive field maintenance

In the lateral geniculate nucleus (LGN) and SC, neurons undergo critical periods of plasticity during which visual experience is required for the maintenance of previously refined RFs (Carrasco et al., 2005) in SC, and to prevent exuberant innervation from retinal ganglion cell (RGC) afferents (Hooks & Chen, 2006; Hooks & Chen). Imposing a wavelength that is well matched to the size of refined RFs in the LGN and SC is one way in which visual experience may prevent the loss of refinement that is seen in visually deprived animals. In the primary visual cortex, neurons exhibit immature RFs at the onset of vision, but refine in size and selectivity for visual features such as the orientation of contours, and the direction of motion, with visual experience (Hubel & Wiesel, 1963; Braastad & Heggelund, 1985; Chapman & Stryker, 1993; White et al., 2001; Li et al., 2006 ; Li et al., 2008). The patterned structure of the visual environment can influence this process. For example, restricting experience to a subset of orientations promotes the maturation of RFs that respond to only those exposed features (Stryker et al., 1978). Orientation selectivity is thought to arise from the alignment of inputs from the LGN along an axis parallel to the preferred orientation (Hubel & Wiesel, 1962), and their spatial frequency tuning arising from the periodic arrangement of these inputs (Movshon

et al., 1978). Our results offer one mechanism for how this arrangement of inputs from the LGN may arise in cortical simple cells (Figs. 4.2 and 4.4). This raises the question of whether natural visual environments contain the appropriate spatiotemporal properties for refinement by the mechanism we describe, as the wave speed required for refinement is constrained by the target size and shape of the RF. Efforts to model the spatiotemporal structure of natural scenes (Dong & Atick, 1995; Reinagel & Zador, 1999) provide some insight, but are limited in their application to particular cases of visual experience, making it difficult to extrapolate how visual contours may move through the visual environment to elicit travelling wave activity patterns in the retina. Nevertheless, we have shown in Fig. 4.3 that a range of wave speeds and STDP properties are able to refine a RF, such that variability in moving contours may not pose too great a problem for a particular type of RF refinement, whether it involves division into subfields or a change in size. Indeed, we showed that even noisy wave patterns with complex structures were able to refine RFs in our simulations (Fig. 4.6). Implicit in this strategy is the assumption that retinal ganglion cells (RGCs) generate wave like activity patterns when exposed to moving natural scenes. Given common properties of RGCs, which include sparse firing and transient spike bursts to local changes in luminance (Pitkow & Meister, 2012; Carandini et al., 2005), in addition to travelling waves through RGC populations recorded *in vitro* in response to moving light stimuli (Frechette et al., 2005), it seems likely that motion in the visual environment would elicit wave-like patterns of activity in the retina.

Physiological changes during development that may aid refinement

Our results also suggest a potential role for physiological changes that occur alongside the anatomical refinement of neural circuits during early developmental stages, by way of their impact upon wave propagation and STDP. For example, stage 2 retinal waves in mice, which exert a strong influence over the refinement of RFs

in the SC, gradually decrease in speed by approximately 30% over a period of 5–6 days (Maccione et al., 2014). Decreases in STDP time constants, on the other hand, may result from developmental changes at individual synapses and changes in cell membrane conductances. The composition of NMDA receptors (NMDARs) has been associated with the decay time of STDP rules, with long decay times being associated with the prevalence of NMDARs containing NR2B subunits, and short decay times being associated with NMDARs containing NR2A subunits (Guo et al., 2012). The increased expression of NR2A relative to NR2B containing NMDARs throughout the brain during early postnatal development (Paoletti et al., 2013) may therefore provide a developmental decrease in STDP time constants. Increases in postsynaptic conductance also reduce the decay time of STDP rules (Fuenzalida et al., 2007; Delgado et al., 2010; Hao & Oertner, 2012), an effect which may be enhanced by developmental increases in neuronal resting membrane conductances (McCormick & Prince, 1987; Etherington & Williams, 2011). In Fig. 4.4, we demonstrated how both the wave speed and the decay time constants of STDP determine the spatial frequency that is imposed on developing RFs, thereby determining how a RF is refined. Consequently, continual changes to these properties over the course of development would enable the continual refinement of a RF. Our results therefore suggest one possible utility for changes in wave speeds and STDP time scales over the course of the development by establishing a relationship between pattern formation and RF refinement, and their dependence on wave and STDP properties.

5. Results

Analysis of candidate STDP rules for retinocollicular refinement and their influence on the role played by retinal waves

Background

Before the onset of vision, spontaneous wavefronts of spiking activity traverse the developing retina in three well defined stages (Torborg & Feller, 2005). The second of these stages is characterised by waves that sweep out large areas at a slow speed, typically 0.1–0.3 mm/s. Disrupting wave activity at this stage, without abolishing retinal ganglion cell (RGC) spiking altogether, has a profound effect on the precise targeting of RGC axons in their recipient nuclei. In wild type mice, axons emanating from a small area of the retina terminate in small areas of the superior colliculus (SC), whereas RGC axons exhibit large termination zones over broad areas of the SC in mice that have perturbed stage 2 retinal waves (McLaughlin et al., 2003a; Chan-

drasekaran et al., 2005; Pfeifferberger et al., 2006; Huberman et al., 2008a; Dhande et al., 2011). Conversely, when retinal waves are intact, but instead the guidance of RGC axons by molecular cues is disrupted, axons emanating from a small area of the retina terminate in multiple ectopic, yet small, areas of the SC (Pfeifferberger et al., 2006; Chandrasekaran et al., 2005; Chandrasekaran et al., 2009). In accordance with these anatomical data, individual SC neurons exhibit considerably enlarged receptive fields (RFs) when stage 2 retinal waves are perturbed (Chandrasekaran et al., 2005) and patchy RFs when guidance cues are perturbed (Chandrasekaran et al., 2009). Taken together, these phenomena prompted the hypothesis that the spatiotemporal correlations of retinal waves are instructive for the early development of retinocollicular circuits through Hebbian mechanisms of synaptic plasticity. Butts & Rokhsar (2001) calculated that, during retinal waves, the burst times of RGCs conveyed the most information about their relative retinotopic positions over a time scale of 0.1–2 s, suggesting that mechanisms involving timing dependent plasticity would best make use of this information for the refinement of RGC axons if it were sensitive to this time scale.

Our findings in Chapters 3 and 4 agree with the Butts & Rokhsar (2001) hypothesis. When applied to retinal wave dependent development of the visual system, our model predicted that STDP rules with time scales of tens of milliseconds would be unsuitable for the periodic patterning of feedforward circuits, and thus the refinement of RFs, for two reasons: 1) the slow speeds characterising retinal waves require STDP rules that operate over time scales of 0.2–1 s in order for RFs to be refined to sizes that are observed experimentally; 2) retinal waves are characterised by burst durations > 1 s, such that changes in the connectivity pattern that are due to the distinct spatiotemporal correlations induced by propagating wavefronts would be occluded by changes due to uncorrelated activity within single spike bursts, thus preventing the development of periodic connectivity patterns in feedforward circuits

(Fig.s 3.3C and 3.4 in Chapter 3).

If the precise spatiotemporal correlations induced by retinal waves are important, insofar as they can instruct refinement through periodic patterning, the reasons outlined above suggest a need for STDP rules with long decay time constants. However, this requirement challenges a large fraction of the existing data for STDP rules, which do not exhibit the long decay time constants predicted by our model (Caporale & Dan, 2008). Furthermore, several studies that utilised more naturalistic spike trains to investigate STDP, rather than single pre- and postsynaptic spike pairs, revealed that synaptic plasticity cannot be fully explained by the linear addition of synaptic updates by an STDP rule (Froemke & Dan, 2002; Froemke et al., 2006). Rather, the magnitude and direction of synaptic plasticity is sensitive to a number of factors including the firing rate, postsynaptic membrane voltages, postsynaptic intracellular calcium concentrations, the current synaptic strength, neuromodulation and the dendritic location of the synapse (Markram et al., 1997; Ngezahayo et al., 2000; Sjöström et al., 2001; Froemke et al., 2005; O'Connor et al., 2005; Nevian & Sakmann, 2006; Sjöström & Häusser, 2006; Seol et al., 2007).

These results have motivated more sophisticated models of synaptic plasticity that go beyond the phenomenological model of STDP, like that of Song et al. (2000) and Froemke & Dan (2002), to more mechanistic models that depend upon the postsynaptic voltage or intracellular calcium concentration, $[Ca^{2+}]$ (Clopath et al., 2010; Graupner & Brunel, 2012). The Graupner & Brunel model recapitulates several phenomenological STDP rules for single spike pairs with a time scale determined, in part, by the decay time constant of $[Ca^{2+}]$, which is itself largely dependent on the decay time of N-methyl-D-aspartate receptors (NMDARs) (Sabatini et al., 2002). However, the vast majority of studies that report short time scales for STDP in mice are conducted after the second postnatal week, whereas stage 2 retinal waves

occur during the first postnatal week. At around P10, the NMDAR subunit composition in mice undergoes a switch from a dominance of NR2B subunits, which have long decay times, to a dominance of NR2A subunits, which have much shorter decay times (Liu et al., 2004; Liu & Chen, 2008; Paoletti et al., 2013). Thus, stage 2 retinal waves and retinofugal refinement are active at a time when STDP rules are likely to have much longer time scales than those often observed in older animals.

Recent studies support this idea. In the cortex of 3–4 week old dark reared mice, STDP rules with time scales that extend beyond 100 ms were observed, whereas the time scale for STDP in age-matched normally reared mice did not exceed 50 ms (Guo et al., 2012). Guo et al. argue that this is due to an increased ratio of NR2B:NR2A subunits in NMDARs, brought about by dark rearing. In another study at the retinogeniculate synapses of rats in their first postnatal week, pairing pre- and post-synaptic bursts 1 s in duration, which are characteristic of retinal waves, induces synaptic plasticity in a manner that depends on the relative timing of the bursts, rather than individual spikes (Butts et al., 2007). The window for LTP in this burst-timing dependent plasticity (BTDP) rule was ± 0.5 s, with LTD occurring for burst latencies $0.5 \text{ s} < \Delta t_{burst} < 1.0 \text{ s}$. The exact mechanism underlying BTDP is not yet understood. Butts et al. (2007) suggest one scheme in which BTDP can arise from STDP, whereby synapses are strengthened by spikes that are coincident within 50 ms of each other and are weakened by a constant term, although no biological mechanism was suggested for weakening.

At present, altogether little is known about STDP at retinocollicular synapses during stage 2 retinal waves. We therefore explore two candidate models of STDP that are able to reproduce the BTDP rule described by Butts et al. (2007), and examine and contrast their influence on the development of RFs. First, we show that the traditional implementation of STDP rules with long time scales can give rise

to BTDP rules, and that such rules can build periodic connectivity patterns and refine RFs when driven by characteristic retinal wave activity in the same manner as shown in Chapters 3 and 4. Second, we examine a more mechanistic model of STDP, the $[Ca^{2+}]$ model of Graupner & Brunel (2012), and demonstrate the equivalence between this model and the STDP model proposed by Butts et al. (2007), when suitable parameters are chosen, in their ability to yield the BTDP rule. Using our own implementation of the Butts et al. STDP model as a proxy for the $[Ca^{2+}]$ model, we show that it does not lead to periodic patterning, but demonstrate that it can nevertheless lead to the refinement of RFs in a manner that does not require travelling wave input patterns. Finally, we demonstrate that the Butts et al. STDP model and traditional STDP rules with long time scales have opposite effects on the development of RFs when the propagation of retinal waves is disrupted. This motivates discussion of recent studies in which the healthy refinement of retinotopy in the mouse SC was obtained even after perturbed retinal wave propagation (Xu et al., 2011; Zhang et al., 2012).

Methods

Retinal wave simulations

To generate complex and noisy waves in 2D, we implemented a simulation based on a previously published model of the developing mammalian retina (Feller et al., 1997). A full description of the retina model is provided in General Methods. Here, we describe specific modifications made for the simulations in this chapter.

In order to confirm that waves generated by this model instruct RF development in the same speed dependent manner as was described in Chapter 4, we scaled the spike times generated by the simulation by a constant factor, F_t , in order to obtain

retinal wave simulations with faster and slower wave speed distributions. Scaling time would not only change the speed of wave propagation, but would also change the length of bursts. However, we showed in Chapter 3 that the burst duration could significantly affect the outcome of patterning in the network’s connectivity. To hold the burst duration constant, we isolated every burst of spikes generated by the retinal wave model for each neuron and multiplied the burst onset time by $F_t = 1.4$ to produce slow waves, 1.0 for medium speed waves and 0.6 for fast waves. New spike times for spikes within that burst, t_i^{new} , were then calculated from the original spike times, t_i^{old} , using $t_i^{\text{new}} = t_i^{\text{old}} - (1 - F_t)t_{\text{onset}}^{\text{old}}$, where $t_{\text{onset}}^{\text{old}}$ was the original burst onset time.

Input activity patterns of stationary spots

Input patterns consisting of stationary spots of activity were generated in order to test the necessity for the distinct spatiotemporal correlations of travelling waves. Each spot was initiated at a random location on the input layer with a given diameter. All input neurons within this diameter generated spikes using a Bernoulli process for a burst duration of 1 s. Immediately following the termination of one spot, a new spot was initiated.

Ca²⁺ dependent model of synaptic plasticity

The common principal underlying several computational models of calcium dependent plasticity is that synaptic strengthening and weakening are proportional to a non-linear function of the postsynaptic intracellular calcium concentration, $[\text{Ca}^{2+}]$, whereby low $[\text{Ca}^{2+}]$ does not lead to any change in synaptic strength, moderate $[\text{Ca}^{2+}]$ leads to synaptic weakening, and high $[\text{Ca}^{2+}]$ leads to synaptic strengthening. We employ a relatively simple implementation of this model that was described by Graupner & Brunel (2012), which we hereafter refer to as the $[\text{Ca}^{2+}]$ model.

Synaptic strength dynamics in the full $[\text{Ca}^{2+}]$ model are dependent on the instantaneous strength of the synapse. However, this is not a feature in other models of Ca^{2+} dependent synaptic plasticity (for example, Shouval et al., 2002), and the Graupner & Brunel $[\text{Ca}^{2+}]$ model is designed such that this dependence is negligible during input and output activity. Because we are principally concerned with synaptic dynamics only when the input and output neurons are active, we use the following reduced version of the model:

$$\frac{1}{\eta} \frac{dw(t)}{dt} = \gamma_p H(c(t) - \theta_p) - \gamma_d H(c(t) - \theta_d), \quad (5.1)$$

where $w(t)$ is the synaptic strength and $c(t)$ the intracellular calcium concentration in the output neuron at time, t . A threshold for plasticity is provided by $H(x)$, the Heaviside step function, such that the synapses strengthens at a rate $\eta(\gamma_p - \gamma_d)$ if $c(t) \geq \theta_p$ and weakens at a rate $\eta\gamma_d$ if $c(t) \geq \theta_d$.

The dynamics of the postsynaptic calcium concentration at a synapse is governed to a large extent by calcium influx through NMDA receptors and voltage dependent calcium channels (Sabatini et al., 2001). Calcium influx is therefore modelled by an instantaneous increase in $c(t)$ by an amount C_{in} and C_{out} after each input and output spike, respectively, after which $c(t)$ decays exponentially with a time constant of τ_{Ca} . We follow the example of Graupner & Brunel (2012) and use parameter values that produce a typical, asymmetric STDP rule for single pairs of pre- and postsynaptic spikes. These give $\gamma_p = 321.808$, $\gamma_d = 200$, $\theta_d = 1$, $\theta_p = 1.3$, $C_{\text{in}} = 1$, $C_{\text{out}} = 2$ and $\tau_{\text{Ca}} = 20$ ms. These parameters failed to reproduce the synaptic weakening observed by Butts et al. (2007) when the $[\text{Ca}^{2+}]$ model was applied to unpaired input and output bursts. We therefore tested an adjusted version of the model by setting $\theta_d = 1.7$ and $\theta_p = 1.94$, which succeeded in capturing the time dependence

of BTDP in Butts et al., (2007).

STDP_{sub} model as an implementation of the BTDP rule

In order to produce a temporally symmetric BTDP rule from a STDP that operates on bursts of action potentials, Butts et al. (2007) proposed a modification to the traditional, asymmetric STDP rule. They model the BTDP rule by describing the percentage change in synaptic strength, Δw , for a set of N pre- and postsynaptic spike pairs with latencies that are within 50 ms, using $\Delta w = 0.145N - 4.0472\%$. Thus the first term in the right-hand side describes an STDP component to plasticity that operates linearly on all pre- and postsynaptic spike pairs. The STDP component only gives rise to synaptic strengthening and depends only on the coincidence of pre- and postsynaptic spikes within a 50 ms window, but not their temporal order. The second term provides a source of synaptic weakening. This is modelled by a simple subtraction from the synaptic strength, although the authors do not suggest any biological mechanisms that may implement this subtraction.

In order to translate the Butts et al. BTDP rule into our own simulation framework, we implement the following rule. For a single pair of input and output spikes, separated in time by Δt , synaptic strengthening follows the same procedure as for the STDP rules described in Chapter 3, and is described by a *boxcar* function

$$K(\Delta t) = H(\Delta t + \tau_+) - H(\Delta t - \tau_+), \quad (5.2)$$

where $H(t)$ is the Heaviside step function. As such, $K(\Delta t)$ is centred at $\Delta t = 0$ and as a width of $2\tau_+$. We implement the subtractive component of synaptic weakening by subtracting a constant, C_- , from the synaptic strength after every output

spike. Thus, synaptic weakening is dependent on the output firing rate. This might represent, for example, the release of endocannabinoids from the postsynaptic site after Ca^{2+} influx due to heightened postsynaptic depolarisation (Heifets & Castillo, 2009). Subtractive weakening and a dependence of weakening on output firing rates have both been shown to be useful strategies that help establish structure in developing connectivity patterns (Oja, 1982; Miller & MacKay, 1994). We refer to our implementation of the Butts et al. BTDP rule as the STDP_{sub} model, where ‘sub’ denotes the subtraction term.

Constructing a BTDP rule from the $[\text{Ca}^{2+}]$ and STDP_{sub} models

We computed effective BTDP rules by simulating the experimental protocol used in Butts et al. (2007). Using a Bernoulli process with a mean firing rate of 10 Hz, we generated 1 s input and output bursts for which the burst onset times were offset by an amount $\Delta t_{\text{burst}} = t_{\text{in}} - t_{\text{out}}$, where t_{in} and t_{out} denote the onset times of the input and output bursts, respectively. We varied Δt_{burst} in the range $[-2 \text{ s}, +2 \text{ s}]$ in steps of 40 ms. We fed the same spike sequence for each pair of input and output bursts through both the $[\text{Ca}^{2+}]$ model and the STDP_{sub} model in order to calculate the change in synaptic strength for each value of Δt_{burst} , thus yielding an effective BTDP rule for each model. The stochastic nature of the spike sequences made the effective BTDP rule very noisy. We therefore repeated this process for 100 trials, generating new spike sequences for each trial, and computed the average effective BTDP rule.

Results

Interaction of retinal waves with long time scale STDP rules

In the following, we frame our investigation in the context of RF refinement, not only because of the functional importance of RFs, but also as a description of retinotopic refinement, under the assumption that coarse grained retinotopic information is established and maintained by some other mechanism, such as signalling between guidance molecules. Our findings in Chapters 3 and 4 indicated that the interaction between slow retinal waves and STPD required STDP rules with time constants of the order of hundreds of milliseconds, in order for RFs to be refined to the sizes observed *in vivo* in young mammals. It is not yet clear whether STDP at retinofugal synapses operates over long time scales such as this during stage 2 retinal waves, when considerable refinement takes place. Thus, our goal here is to examine how such STDP rules might impact the development of immature RFs that are driven by retinal waves, so as to determine whether they might be biologically feasible and to later compare them with other models of plasticity.

First, we must address one of the results from Chapter 3, in which we showed that long bursts, greater than 0.5 s in duration, were ineffective in building periodic connectivity patterns under STDP rules with short time scales of 20 ms (Fig. 3.4). The retinal waves that we model in this chapter have burst durations of 1 s and, in experimental recordings, bursts can last up to two or three seconds. However, this may not pose a problem when the STDP rule has a long time scale. This is because STDP rules with short time scales have a large bandwidth in frequency space. As such, they filter out little of the noise due to uncorrelated spikes within a burst. This noise component becomes more dominant as the burst duration increases, and occludes the structure of $\tilde{\kappa}(k)$, the function that determines the dominant spatial

frequency that is imposed on the connectivity pattern by the wave (see Eqs. 3.6–3.9). However, STDP rules with long time scales concentrate power in $\tilde{\kappa}(k)$ over a smaller bandwidth, and therefore do filter out much of the noise due to uncorrelated spikes within a burst. Because refinement of RFs by retinal waves requires values of τ_+ that are in the range 0.2–1.0 s (see Fig. 3.3), an order of magnitude longer than typically observed STDP rules, we expected the interaction of temporally wide STDP rules with long bursts in travelling waves to produce periodic connectivity patterns. We confirmed this by testing whether periodic connectivity patterns developed in our simulation with spiking neurons. We ran simulations of 1D waves, as we had done for Fig. 3.2 in Chapter 3, using a burst duration of 1 s, wave speeds that are reported for retinal waves and STDP time constants predicted by our model, for both symmetric and asymmetric STDP rules. Spikes were generated at a rate of 10 Hz, which matches experimentally observed firing rates (Stafford et al., 2009). In Fig. 5.1A, we provide examples of connectivity pattern development for both the asymmetric and symmetric STDP rules with $v = 0.2$ mm/s and $\tau_+ = 0.4$ s. In Fig. 5.1B, we plot the mean spatial frequencies of stabilised connectivity patterns for a range of STDP time scales ($\tau_+ = 0.2, 0.4, 0.6, 0.8$ and 1.0 s), keeping $v = 0.2$ mm/s constant (top row), and the spatial frequencies over a range of wave speeds ($v = 0.1, 0.2, 0.3, 0.4$ and 0.5 mm/s), keeping $\tau_+ = 0.3$ s constant (bottom row). The spatial frequencies obtained in this set of simulations closely matched our theoretical predictions (grey curves in Fig. 5.1B), which were computed numerically from the Fourier transform of $\kappa(x)$ in Eq. 3.7, with coefficients of determination $R^2 > 0.9$ for each plot Fig. 5.1B.

Experiments in rat retinogeniculate synapses, at the stage when retinal waves are prevalent, show that pairing of pre- and postsynaptic bursts induce changes in synaptic strength that depend on the temporal overlap of the bursts, but do not depend on the temporal ordering of the two bursts (Butts et al., 2007). In other words,




Figure 5.1: **Periodic connectivity patterns are stable solutions to the synaptic strength dynamics for STDP rules with long time scales.** **A.** Development of the connectivity pattern, initialised with a homogeneous distribution of synapses with strength 0.5, from 500 input neurons to a single output neuron, under the influence of an asymmetric (top) and symmetric (bottom) STDP rule. In both plots, $\tau_+ = 0.4$ s and $v = 0.2$ mm/s. **B.** Log-log plots for the spatial frequency of the stabilised connectivity pattern as a function of τ_+ (top) and v (bottom) for both symmetric (left) and asymmetric (right) STDP rules. Grey line: predicted k^* . Markers and error bars: mean spatial frequency of the final connectivity pattern in the spiking simulation $\pm$ SEM. **C.** Fig. 3 from Butts et al. (2007), showing the change in EPSC size after pairing stimulated 1 s bursts in the optic tract with stimulated 1 s bursts in LGN. BTDP rules are shown as solid and dashed lines that are fit to the data. The constant subtraction in the Butts et al. BTDP rule is illustrated by the horizontal, dashed line. **D.** Effective BTDP rules that were computed by applying asymmetric (top) and symmetric (bottom) STDP rules with long time scales to pairs of input and output bursts. The curves shown are the mean synaptic strength changes, averaged over 100 paired bursts. Bursts lasted 1 s and had a firing rate of 10 Hz. Grey: rules using $\tau_+ = 0.1$ s. Black: rules using $\tau_+ = 0.3$ s.

Butts et al. showed that plasticity at this synapse can be described by a BTDP rule that is symmetric in time (Fig. 5.1C, taken from Fig. 3 in Butts et al. (2007)). We therefore require that the STDP rules used in our model also produce this symmetry in time for paired pre- and postsynaptic bursts. Because each input-output spike pair in our model contributes linearly to the change in synaptic strength, we can easily compute an effective BTDP rule by applying the STDP rules to hypothetical spike bursts in the following manner. A burst of input spikes, $S_{\text{in}}(t)$, 1 s in duration, was generated by a Bernoulli process with a mean firing rate of 10 Hz and temporal resolution of 1 ms. The convolution of this spike train with the STDP rule, $K_{\text{in}}(t) = \int_{-\infty}^{\infty} d\Delta t K(\Delta t) S_{\text{in}}(t - \Delta t)$, informs us of the change in synaptic strength if a single output spike were paired with the input burst at time, t . Thus, the total change in synaptic strength due to pairing an output burst, $S_{\text{out}}(t)$, with the input burst, is the inner product, $\Delta w = \int_{-\infty}^{\infty} dt K_{\text{in}}(t) S_{\text{out}}(t)$. We can therefore construct the effective BTDP rule, $K_{\text{BTDP}}(\Delta t)$, as a function of the time lag between a given input and output burst, by computing the cross-correlogram between $K_{\text{in}}(t)$ and $S_{\text{out}}(t)$, $K_{\text{BTDP}}(\Delta t) = \int_{-\infty}^{\infty} dt K_{\text{in}}(t + \Delta t) S_{\text{out}}(t)$, where $\Delta t = t_{\text{in}} - t_{\text{out}}$. We computed effective BTDP rules using both asymmetric and symmetric STDP rules, and plot the effective BTDP rules averaged over 100 input-output burst pairs in Fig. 5.1D. Because the effective BTDP rule is constructed by a combination of linear operations, the asymmetric STDP rule necessarily produced an asymmetric BTDP rule (top panel), and therefore is not a good candidate for BTDP at the retinogeniculate synapse using this approach. For the same reason, the symmetric STDP rule is a good candidate for BTDP. However, when $\tau_+ = 0.1$ s (grey curve, bottom panel), synaptic weakening dominates and very little strengthening is obtained for overlapping bursts. A longer decay time of $\tau_+ = 0.3$ s (black curve, bottom panel) recovers the main features of BTDP as observed in Butts et al. (2007), with a central positive lobe and flanking negative lobes. Even longer values for τ_+ must also reproduce these main features.

Refinement of receptive fields by retinal waves and STDP with long decay times

Having shown that the mechanism for periodic pattern formation in feedforward circuits, laid out in Chapter 3, also applies to 1D retinal waves and long STDP time scales, we next confirmed that 1D retinal waves can also give rise to the refinement of RFs. We ran simulations that were initialised with a RF, setting $RF_0 = 0.8$ mm, which is an estimate of the retinal projection of a stimulus that matches the size of an unrefined RF in the rodent SC (Chandrasekaran et al., 2005), in the same manner as for Fig. 4.2. The development of the RF in the simulation is shown in Fig. 5.2A, using $v = 0.2$ mm/s and $\tau_+ = 0.7$ s.

In Chapter 4, we showed that complex, fast travelling wave patterns were able to refine RFs with an immature size. In order to draw comparisons between STDP with long time scales and the Butts et al. BTDP rule later in this chapter, we first tested whether similar, complex wave patterns, which moved slowly and generated long lasting bursts, would also refine RFs in a manner dependent on τ_+ and v . Following a similar procedure as in Chapter 4, we used the Feller et al. (1997) retinal wave model to generate complex wave patterns and scaled these in time by a factor, F_t , to produce different distributions of wave speeds that match those of retinal waves, while maintaining the gross spatiotemporal structure of the wave patterns. We took additional steps to ensure that bursts were ~ 1 s in duration with a firing rate of 10 Hz (see Methods; Wong et al., 1993; Stafford et al., 2009). We ran simulations that used these wave patterns as inputs to a single output neuron, which was set up with an initial circular RF ($RF_0 = 0.82$ mm).

In the 1D simulations, we used $\tau_+ = 0.7$ s to achieve refinement. However, we know from Chapter 4 that 2D waves, which travel in multiple directions, require $\kappa(x)$ to comprise higher spatial frequencies than in 1D in order to achieve refine-




Figure 5.2: **Development of RFs using STDP rules with long time scales.** **A.** Example of RF development from a 1D simulation, using $\tau_+ = 0.7$ s and $v = 0.2$ mm/s. Green lines denote boundaries to the arbour function. **B.** Examples of final RFs that developed in simulations under the influence of complex retinal wave patterns that came from different wave speed distributions, and STDP rules that were constructed with different time scales. Centre: RF that was exposed to medium wave speeds; $\tau_+ = 0.3$ s. Top and bottom: RFs that were exposed to slow and fast wave speeds, respectively; $\tau_+ = 0.3$ s. Left and right: $\tau_+ = 0.25$ s and $\tau_+ = 0.35$ s, respectively; RFs were exposed to medium wave speeds. Solid green circles: arbour function boundaries. Dashed green circles: boundaries to the initial RF. Box colours correspond to the RF area plots in **C**. **C.** Development of RF areas (measured as total input area providing synapses with strength ≥ 0.5). Each curve is an average over 3 simulations, using different sequences of retinal wave patterns. Left: RF areas for slow, medium and fast wave speeds (dark, medium and bright blue, respectively); $\tau_+ = 0.3$ s. Right: RF areas for $\tau_+ = 0.25$ s, $\tau_+ = 0.30$ s and $\tau_+ = 0.35$ s (dark, medium and bright purple, respectively); medium wave speeds. Dashed grey lines: initial RF size.

ment, as RFs tend to expand in the direction orthogonal to wave propagation. We therefore ran simulations using $\tau_+ = 0.3\text{ s}$ and set F_t such that the median wave speed was 0.23 mm/s (calculated using the sCCG, see General Methods). The centre panel in Fig. 5.2B provides an example of a RF that had refined to a smaller area and maintained an approximately circular profile under these conditions. We next varied F_t and τ_+ to confirm that the characteristic spatial frequency of $\kappa(x)$ was underlying the refinement of the RF in these simulations, and provide examples of the final RFs for each set of parameters in Fig. 5.2B. Using the same wave inputs (median wave speed of 0.23 mm/s), we set $\tau_+ = 0.25\text{ s}$ and 0.35 s (Fig. 5.2B, left and right panels, respectively). Next, we kept $\tau_+ = 0.30\text{ s}$ constant for simulations using different F_t values, such that the wave speed distribution had a median of either 0.17 mm/s or 0.35 mm/s (Fig. 5.2B, top and bottom panels, respectively). These simulations confirmed that larger RFs developed for faster wave speeds (Fig. 5.2C, left panel) and longer values of τ_+ (Fig. 5.2B, right panel), as predicted by our theoretical results in Chapter 3.

STDP rules with short decay times that yield burst-timing dependent plasticity

Since the detailed characterisation of STDP by Zhang et al. (1998) and linear implementation of STDP rules in computational models by Song et al. (2000), a number of experiments have shown that changes in synaptic strength cannot be fully explained by the linear contribution of individual spike pairs acting through a STDP rule (Sjöström et al., 2001; Froemke et al., 2005; Nevian & Sakmann, 2006). Of several more biophysical, non-linear models of STDP, models in which changes in synaptic strength depend on intracellular calcium concentration, $[\text{Ca}^{2+}]$, have received considerable attention (Shouval et al., 2002; Graupner & Brunel, 2012). In these models, Ca^{2+} enters the postsynaptic cell in response to pre- and postsynaptic spikes, for example, via NMDARs and voltage dependent calcium channels (VD-

CCs), which leads to an elevation in $[\text{Ca}^{2+}]$. In general, small $[\text{Ca}^{2+}]$ levels do not lead to changes in synaptic strength, whereas moderate $[\text{Ca}^{2+}]$ induces LTD, and large levels of $[\text{Ca}^{2+}]$ lead to LTP. Graupner & Brunel. (2012) introduced a simplified model of $[\text{Ca}^{2+}]$ dependent plasticity (described in Methods), with which they were able to reproduce results from a variety of different plasticity protocols that generated STDP rules with different shapes. We adopt their model in this section, using parameters that produce the classical, asymmetric STDP rule, and demonstrate its equivalence to the model of BTDP proposed by Butts et al. (2007), which we later implement in simulations to examine its influence on the development of connectivity patterns.

In the Graupner & Brunel (2012) model, impulse responses for $[\text{Ca}^{2+}]$, after a spike from an input or output neuron, decay exponentially with a time constant of 20 ms, and are drawn in Fig. 5.3A (red: input response; blue: output response). Synapses are weakened by an amount proportional to the time (marked by grey arrows in Fig. 5.3A) spent by $[\text{Ca}^{2+}]$ above the threshold for LTD, θ_d , with γ_d the constant of proportionality. Likewise, synaptic strengthening is proportional to the time (marked by black arrows in Fig. 5.3A) spent above the threshold for LTP, θ_p , with γ_p the constant of proportionality. The model is set up such that $[\text{Ca}^{2+}]$ does not cross either threshold after single, unpaired input spikes, whereas it crosses both thresholds after single, unpaired output spikes. One can set the ratio γ_d/γ_p such that unpaired output spikes lead to equal strengthening and weakening. Thus, when the difference in spike times, $\Delta t_{\text{spike}} = t_{\text{in}} - t_{\text{out}}$, is large, synaptic strengths remain unchanged. In Fig. 5.3B (purple), we draw the effective STDP rule that is obtained using the Ca^{2+} model for a single input-output spike pair as a function of Δt_{spike} , which strongly resembles the typical asymmetric STDP rule observed in experiments. However, it is difficult to analyse the Graupner & Brunel model and its interaction with travelling waves using the same approach as in Chapter 3, as the phenomenological rule that

is obtained by pairing input and output spikes is very sensitive to the presence of multiple spikes. We therefore explore another rule, the STDP_{sub} rule, which we will show to behave in a similar manner as the $[\text{Ca}^{2+}]$ model, but that is more tractable.

Butts et al. (2007) also proposed an alternative to the traditional STDP rule, in order to reproduce the BTDP rule when applied to paired 1 s bursts of pre- and postsynaptic spikes. Their BTDP rule comprises two components. The first component operates in the same manner as a traditional STDP rule, whereby synapses are strengthened if input and output spikes occur within 50 ms of each, regardless of their temporal order. In the second component, a constant is subtracted from this STDP rule to provide synaptic weakening. However, there is not a unique way in which this second component could be implemented when there are multiple pairs of bursts, as is the case in our simulations of connectivity development. Therefore, in order to implement the Butts et al. BTDP rule in our own framework, we devised a similar model, STDP_{sub} , which emulates the two components in the BTDP rule (see Methods). The implementation of synaptic strengthening by STDP follows the same procedure as the STDP rules used in Chapters 3 and 4, and exactly matches strengthening in the BTDP rule described by Butts et al. However, our implementation of synaptic weakening differed slightly, whereby all synapses that were afferent to the output neuron were decremented by an amount C_- following every output spike. This provided the necessary synaptic weakening for unpaired bursts, which is a property of the Butts et al. BTDP rule. As such, synapses in the network are suppressed by an amount proportional to the output firing rate. This bears analogy to subtractive normalisation and the rate dependent Oja rule, which are well known means of building structured connectivity patterns with Hebbian plasticity (Oja, 1982; Miller & MacKay, 1994).

In the following, we apply the $[\text{Ca}^{2+}]$ model and STDP_{sub} model to paired bursts

of input and output spikes, each 1 s in duration with a mean firing rate of 10 Hz, in order to compute effective BTDP rules for each model (see Methods). Input and output bursts are paired over a range of temporal offsets, $\Delta t_{\text{burst}} = t_{\text{in}} - t_{\text{out}}$, where t_{in} and t_{out} now denote the onset time of each burst. For each Δt_{burst} , new input and output bursts are generated by a Bernoulli process and the change in synaptic strength computed in order to yield an effective BTDP rule. This process is repeated over 100 trials so that an average BTDP rule can be calculated.

In the $[\text{Ca}^{2+}]$ model, the time spent by $[\text{Ca}^{2+}]$ above θ_p is determined by the maximum $[\text{Ca}^{2+}]$, which will be greater with more coincident spikes. Likewise, in the STDP_{sub} model, the total number of strengthening synaptic updates will be proportional to the number of coincident spikes. Therefore, for both models, we expect the change in synaptic strength to be proportional to the temporal overlap of the two bursts. In Fig. 5.3C, we plot the effective BTDP rules obtained using both the STDP_{sub} model (left panel) and the $[\text{Ca}^{2+}]$ model (centre and right panels). The STDP_{sub} model yields a BTDP rule for which synaptic strengthening occurs when $|\Delta t_{\text{burst}}| < 0.7$ s and weakening otherwise, which closely matches the findings in Butts et al. (2007). However, the $[\text{Ca}^{2+}]$ model yielded a BTDP rule that, averaged over all trials, exhibited only strengthening when using the parameters given in Graupner & Brunel (2012) (Fig. 5.3C, middle). Adjusting the thresholds for plasticity such that $\theta_d = 1.7$ and $\theta_p = 1.94$ aligned the behaviour of the $[\text{Ca}^{2+}]$ model with that of the STDP_{sub} model, whereby synaptic strength changes little after isolated input bursts, weakens after isolated output bursts and strengthens for coincident input and output bursts (Fig. 5.3C, right). Although the $[\text{Ca}^{2+}]$ model with adjusted thresholds is considerably more noisy than the STDP_{sub} model, they both yield near identical BTDP rules on average (Fig. 5.3D). Because of this equivalence, and due to the tractability of the STDP_{sub} model and ease with which it can




Figure 5.3: **Effective BTDP rules using non-linear models of STDP.** **A.** $[Ca^{2+}]$ impulse response to input (red) and output (blue) spikes in the Ca^{2+} model of synaptic plasticity. Thresholds for LTD and LTP are shown by horizontal dashed lines, denoted θ_d and θ_p , respectively. Inset: magnification of the boxed area. The degree of weakening and strengthening is proportional to the time spent above the respective thresholds (delimited by the grey and black arrows, respectively) by the summed input and output sources of $[Ca^{2+}]$. **B.** Effective STDP rules for single pairs of input and output spikes, constructed using the Ca^{2+} model (purple) and the $STDP_{\text{sub}}$ model (green). The constant subtraction is $C_- = 0.3$. **C.** Effective BTDP rules, constructed by applying the non-linear STDP models to pairs of input and output bursts. Input and output bursts were paired for each value of Δt_{burst} tested (in the range $[-2\text{s}, 2\text{s}]$ in steps of 40 ms). This was repeated 100 times using different spike sequences for each burst. Plots of the effective BTDP rules show the mean change in synaptic strength (light green and dark purple) and the SD in synaptic strength changes (dark green and light purple). Left: effective BTDP rule that results from the $STDP_{\text{sub}}$ model, using the same value of C_- as in B. Middle: effective BTDP rule that results from the Ca^{2+} model, using parameters taken from Table S1 in Graupner & Brunel (2012). On average, these parameters fail to produce LTD for unpaired bursts ($|\Delta t_{\text{burst}}| > 1\text{s}$). Right: raising the thresholds for plasticity, by increasing θ_d and θ_p , does yield LTD for unpaired bursts on average. **D.** Comparison of effective BTDP rules demonstrates the equivalence of the Ca^{2+} model and the $STDP_{\text{sub}}$ model when applied to 1 s bursts.

be implemented in our simulations, we use the STDP_{sub} model as a proxy for the $[\text{Ca}^{2+}]$ model in the following sections.

The interaction of travelling waves with the STDP_{sub} model

The interaction of travelling waves with the STDP_{sub} model is amenable to the same analytical treatment that we pursued in Chapter 3 for more typical STDP rules. We can query the type of spatial structure this new plasticity model will impose on feedforward connectivity patterns by computing the Fourier transform of each component in the model. The STDP component, $K(\Delta t)$, is a *boxcar* function centred on $\Delta t = 0$. The real component comprising the Fourier transform of a *boxcar* function is a *sinc* function ($\sin(k)/k$), for which power is greatest at low spatial frequencies and has maximum amplitude at the zeroth spatial frequency. However, the second component to the STDP_{sub} model is a constant subtraction from the synaptic strength that follows every output spike. The Fourier transform of a constant is a delta function, $\delta(k)$, which is zero-valued everywhere except at $k = 0$. Subtraction of a constant in the STDP_{sub} model therefore reduces power at only the zero-frequency of $\tilde{\kappa}(k)$. Note that the dynamics of synaptic strengths under the influence of travelling waves in the inputs, as described by Eq. 3.5 in Chapter 3, also depend on the input burst waveform, $\alpha(t)$, which is also a *boxcar* function, and the excitatory postsynaptic potential, $\epsilon(t)$, which is a difference of exponentials. The Fourier transforms of $\alpha(t)$ and $\epsilon(t)$ also exhibit the greatest power at the lowest spatial frequency. Thus, when the STDP_{sub} model is driven by travelling wave activity, our analytical results suggest that the fastest growing spatial frequency in the connectivity pattern will always correspond to a single cycle that spans the entire input layer, irrespective of the wave speed or time scale for STDP. However, nonlinearities in the simulation, which are imposed to bound the synaptic strengths, might allow multiple spatial frequencies to exhibit substantial power. Nevertheless, we expected low spatial frequencies to dominate the power spectrum of connectivity

patterns that develop with the STDP_{sub} model.

To test our hypothesis, we ran simulations of 1D waves, as we had done for Figs. 5.1A and B, using the STDP_{sub} model. In a set of preliminary simulations, we determined $C_- = 0.0005/\text{spike}$ to provide balanced synaptic weakening and strengthening for the case $v = 0.2\text{ mm/s}$ and $\tau_+ = 50\text{ ms}$, such that approximately half of the synapses were driven to zero strength and half to maximum strength. In the next set of simulations, we followed the approach taken in Chapter 3 and varied the wave speed and time scale of STDP to determine their influence on the connectivity patterns that developed.

First, we varied the wave speed between simulations in the range $0.1\text{--}0.35\text{ mm/s}$, keeping $\tau_+ = 50\text{ ms}$ fixed. Examples of the development of connectivity patterns are provided in Figs. 5.4A and B for $v = 0.2$ and 0.3 mm/s , and the power spectra for all wave speeds, averaged over repeated trials, given in Fig. 5.4E. As expected, low spatial frequencies dominate the power spectrum, with maximum power occurring at spatial frequencies in the range $0.24\text{--}0.44\text{ cycles/mm}$ across all the wave speeds tested. For comparison, waves with speeds of 0.2 mm/s , and traditional STDP rules with time constants of 50 ms , were predicted to build connectivity patterns with characteristic spatial frequencies in excess of 10 cycles/mm (Figs. 3.3A and B in Chapter 3). However, a trend for higher spectral bandwidths with slower wave speeds was apparent, reflecting the fact that $\tilde{\kappa}(k)$, which is a *sinc* function, must become broader as $\kappa(x)$ narrows with decreasing speed.

An additional feature of the connectivity patterns that developed with the STDP_{sub} model is that the number of strong synapses, as a percentage of all synapses, increased from 13% to 90% with increasing wave speed (Fig. 5.4G, black). This is not related to the fact that more power resides at lower spatial frequencies when

the wave speed is faster, as an equal number of synapses would both strengthen and weaken at all non-zero spatial frequencies. Rather, note that faster waves cast larger wavefronts on the input layer, i.e. they recruit more input neurons at a given time. We therefore hypothesised that the total activity over the input population per unit time is responsible for the increasing number of strong synapses with wave speed. To test this, we repeated these simulations and scaled the firing rate of input neurons, such that if the wave speed was doubled, the firing rate was halved. Fig. 5.4C shows an example of the connectivity pattern development for a wave speed of 0.3 mm/s and input firing rate of 6.7 Hz, for comparison with Fig. 5.4B, in which the firing rate was not scaled. Scaling the input firing rate with the wave speed indeed preserved the number of strong synapses that emerged in the final connectivity patterns (Fig. 5.4G, grey; $p = 0.24$, t -test for the dependence of the number of strong synapses on wave speed with scaled input firing rates), without affecting the dependence of spectral bandwidth on wave speed (Fig. 5.4F). Compare this with the traditional STDP rules used above, for which the wave speed does not exhibit a strong influence over the number of strong synapses that develop. In Fig. 5.4G, we demonstrate the influence of wave speed over the number of strong synapses that develop with the symmetric STDP rule (dark purple: input firing rate was not scaled; light purple: input firing was scaled), for which changes in wave speed between 0.1 mm/s and 0.5 mm/s lead to only a $\sim 3\%$ difference in the number of strong synapses. Although the symmetric STDP rule was sensitive to the input firing rate, such a small change should be expected, as scaling the amplitude of the input firing rate, $\alpha(t)$ in Eq. 3.5 in Chapter 3, simply scales the rate of change in the eigenmodes associated with each spatial frequency. The same trend was found for the asymmetric STDP rule. In summary, with the STDP_{sub} model, the wave speed influences the structure of the connectivity pattern by way of increasing the total input per unit time, which increases the percentage of strong synapses, as well as increasing the spectral bandwidth. In later sections of this chapter, we examine

in more detail the importance of wavefront size on RF refinement.

In a second set of simulations, we varied the time scale for STDP in the STDP_{sub} model in the range 20–70 ms, keeping $v = 0.2 \text{ mm/s}$ fixed. An example of the connectivity pattern development when $\tau_+ = 70 \text{ ms}$ is provided in Fig. 5.4D, for comparison with Fig. 5.4A, and power spectra for all τ_+ values, averaged over repeated trials, given in Fig. 5.4H. Varying τ_+ in this range had little effect on the overall structure of the final connectivity patterns and the spatial frequencies with which they were composed. Note that Eq. 3.7 in Chapter 3 states that $\kappa(x)$ is constructed by a convolution of the STDP rule with additional functions that describe the input firing rate and excitatory postsynaptic potential. In the present case, the long duration of the input burst acts as a low pass filter, suppressing the higher spatial frequencies in $\tilde{\kappa}(k)$ that arise from the short temporal window of the STDP rule. The burst duration therefore presides over τ_+ in determining the spatial frequency components of $\kappa(x)$, when τ_+ is on the order of tens of milliseconds.

The burst duration also governs the wavefront size for a given wave speed, and durations of 1–4 s are typically observed for retinal waves (Stafford et al., 2009; Xu et al., 2011). So as to further affirm the importance of wavefront size with the STDP_{sub} model, we varied the burst duration in another set of simulations, keeping $v = 0.2 \text{ mm/s}$ and $\tau_+ = 50 \text{ ms}$ fixed. Burst durations in the range 0.5–1.75 s were chosen to produce wavefront sizes that matched those in simulations for which the wave speed was varied. Two examples of the development of the connectivity pattern, for burst durations of 1 s and 0.5 s, are given in Figs. 5.5A and B, respectively. When the input firing rate was not scaled with the burst duration, but was fixed at 10 Hz, the number of strong synapses increased from 16% to 87% (Fig. 5.5F, black). This was similar to the increase in the number of strong synapses with increasing wave speed (although the small difference was significant: $p < 0.001$, ANCOVA




Figure 5.4: **Development of the connectivity pattern with the STDP_{sub} model.** **A-D.** Examples for the development of the connectivity pattern and the effect of wave speed (cf. **A** and **B**), the input firing rate (cf. **B** and **C**) and the STDP time scale (cf. **A** and **D**). Values of the free parameters are provided at the top of each panel. **E.** Mean spatial frequency power spectra of the stabilised connectivity patterns in simulations between which the wave speed was varied. Each spectrum is the mean over 16 repeated trials for a given wave speed. Light blue to black: wave speeds from 0.1 mm/s to 0.35 mm/s . **F.** Mean spatial frequency power spectra of the stabilised connectivity patterns in simulations between which both the wave speed and input firing rate was varied, thus maintaining a constant level of input activity per unit time. Light blue to black: wave speeds from 0.1 mm/s to 0.35 mm/s and input firing rates from 20 Hz to 5.7 Hz . **G.** The percentage of synapses with a strength > 0.5 for stabilised connectivity patterns in simulations between which the wave speed was varied. Black and grey: synaptic plasticity was governed by the Butts et al scheme. Black: input firing rate was not scaled with wave speed. Grey: input firing rate was scaled with wave speed. Light and dark purple: synaptic plasticity was governed by the traditional implementation of the symmetric STDP rule. Light purple: input firing rate was not scaled with the wave speed. Dark purple: input firing rate was scaled with the wave speed. The horizontal axes are scaled according to the plasticity rule under consideration: black scale for the Butts et al. scheme and purple scale for the symmetric STDP rule. Markers and error bars show the mean $\pm$ SEM. **H.** Mean spatial frequency power spectra of the stabilised connectivity patterns for simulations between which τ_+ was varied for the STDP_{sub} model. Light blue to black: τ_+ values from 20 ms – 70 ms .




Figure 5.5: **Influence of the burst duration on the development of the connectivity pattern.** **A-C.** Examples of connectivity pattern development. The free parameters of burst duration and input firing rate are shown above each panel. **D.** Mean spatial frequency power spectra of stabilised connectivity patterns for simulations between which the burst duration was varied. Light blue to black: burst durations of 0.5 s to 1.75 s. **E.** Mean spatial frequency power spectra of stabilised connectivity patterns for simulations between which both the burst duration and input firing rate was varied, thus maintaining a constant level of input activity per unit time. Light blue to black: burst durations of 0.5 s to 1.75 s and input firing rates of 20 Hz to 5.7 Hz. **F.** The percentage of synapses with strength > 0.5 in the stabilised connectivity pattern. Black: input firing rate was not scaled with the burst duration. Grey: input firing rate was scaled with the burst duration. Markers and error bars show the mean $\pm$ SEM.

test for covariance of the number of strong synapses with wavefront size), and the spectral bandwidth of the connectivity patterns also increased with burst duration (Fig. 5.5B). However, scaling the input firing rate such that, if the burst duration was doubled, the firing rate was halved, maintained the same percentage of strong synapses for all burst durations tested (Fig. 5.5F, grey; $p = 0.40$, t -test for the dependence of the number of strong synapses on burst duration with scaled input firing rates), without qualitatively affecting the influence of burst duration on the spectral bandwidth (Fig. 5.5E). An example of the development of the connectivity pattern, for a burst duration of 0.5 s and input firing rate scaled to 20 Hz, is given in Fig. 5.5C, for comparison with the unscaled firing rate case in Fig. 5.5B. When the input firing rate was scaled, the percentage of strong synapses at the end of the simulations were not significantly different ($p = 0.99$, two sample t -test) if the wavefront size was varied either by changing the wave speed (Fig. 5.5G) or by changing the burst duration (Fig. 5.4F).

In summary, the behaviour of the STDP_{sub} model differs from traditional STDP rules in two ways. The first difference is that the dominant spatial frequency, k^* , imposed on the connectivity pattern by travelling waves is always the lowest, non-zero frequency, and is independent of the wave speed, v , and the time scale for STDP, τ_+ . In contrast, with traditional STDP rules, k^* follows a logarithmic relationship with v and τ_+ . The second difference is the sensitivity of the connectivity pattern to the total input activity at a given time. With the STDP_{sub} model, increasing the total input per unit time increases the number of strong synapses, whereas the same modification has little effect with traditional STDP rules. These differences have distinct consequences with regard to the development of RFs when using the STDP_{sub} model, which we explore in the next section.

Refinement of receptive fields by the $STDP_{\text{sub}}$ model

We demonstrated in Chapter 4 that the dominant spatial frequency, k^* , of $\kappa(x)$ in Eq. 3.8 in Chapter 3 determined the way in which a RF would refine. For traditional STDP rules, k^* is scaled by the wave speed, v , and the time scale of STDP, τ_+ . However, with the $STDP_{\text{sub}}$ model, k^* is always the lowest possible, non-zero spatial frequency. While the $STDP_{\text{sub}}$ model does not impart a characteristic spatial frequency on the connectivity pattern under travelling wave stimulation, the two components that comprise this scheme compete in such a way as to constrain the number of strong synapses innervating the output neuron. This is a result of global synaptic weakening that occurs whenever the output neuron is active, and strengthening of individual synapses when both the output neuron and respective input neuron are co-active within a STDP window of 50 ms. Consider a connectivity pattern that consists of a large RF. The traversal of a wavefront through the RF drives activity in the output neuron and causes all synapses to be weakened by an amount proportional to the total number of output spikes. However, synapses comprising the edge of the RF will encounter fewer coincident input and output spikes than those comprising the centre of the RF, as the total synaptic input to the output neuron is proportional to the overlap of the wavefront with the RF. As such, synaptic weakening can outweigh strengthening at the RF edges and refine the RF to a smaller size.

To test this, we ran simulations in which the output neuron was initiated with a RF of size $RF_0 = 0.8 \text{ mm}$. In a set of preliminary simulations, we determined a value of $C_- = 0.004$ that would provide such RF refinement for wave speeds of 0.2 mm/s (Fig. 5.6A). Next, keeping C_- fixed, we varied the wave speed between further simulations in the range $0.1\text{--}0.35 \text{ mm/s}$. In Fig. 5.6B, we show the refinement of a RF when $v = 0.1 \text{ mm/s}$. Just as we found above, the slower wave speed forced




Figure 5.6: **Refinement of RFs with the STDP_{sub} model.** **A-C.** Examples of RF development from simulations in which the wave speed and input firing rate were varied. Free parameters are provided above each panel. **D.** The number of synapses with strength > 0.5 as a measure of the stabilised RF size. Black: input firing rate was not scaled with the wave speed. Grey: input firing rate was scaled with the wave speed. Markers and error bars show mean $\pm$ SEM for 16 repeated trials at each wave speed. **E.** Segment of an input activity pattern in which stationary spots appeared in random locations. Each spot was active for a duration of 1 s and its size (delimited by arrows) was equal to the width of a wavefront in simulations that used travelling waves with speeds of 0.2 mm/s as the input. **F-H.** Examples of RF development from simulations that utilised stationary spots as input activity patterns, as shown in **E**. The spot size and input firing rates were varied between simulations. These free parameters are provided above each panel. **I.** The dependence of RF size on spot size, as measured by the number of synapses with strength > 0.5 . Black: input firing rate was not scaled with the spot size. Grey: input firing rate was scaled with the spot size.

more synapses to weaken, further refining the RF. Indeed, the number of strong synapses (having a strength > 0.5) in the final RF, and thus RF size, increased monotonically with wave speed (Fig. 5.6D, black). To confirm this was due to the reduced input per unit time, rather than a reduction in the bandwidth of spatial frequencies imposed on the connectivity pattern, we scaled the input firing rate with the wave speed. As expected, this maintained the number of strong synapses, and therefore also maintained the RF size, across different wave speeds (Fig. 5.6D, grey; $p = 0.94$, t -test for the dependence of the number of strong synapses on wave speed with scaled input firing rates). An example of RF development when $v = 0.1$ mm/s, with the input firing rate scaled to 20 Hz, is provided in 5.6C.

Given that development of the RF depends on wave speed only by way of its influence over the wavefront size, and thus the total input activity at a given time, we suspected that RF development with the STDP_{sub} model does not depend on the travelling nature of the wave. We tested this by generating a new input activity pattern, which consisted of stationary spots of activity that appeared successively at random locations on the input layer. A 10 s excerpt of this activity pattern is provided in Fig. 5.6E. Each spot was active for a duration of 1 s, during which spikes fired at a rate of 10 Hz, immediately after which a spot appeared at a new location. We used this new input activity pattern to drive a single output neuron in a new set of simulations. Spots were equal in size to the wavefront for a given wave speed. Thus, the spot size was varied between simulations in the range 0.1–0.35 mm. In Fig. 5.6F and G, examples of RF development under this new activity pattern are given for spot sizes of 0.2 mm and 0.1 mm, respectively. The final RF size, as measured by the number of strong synapses, increased monotonically with spot size (Fig. 5.6I, black), and closely matched the increase in final RF size with wave speed ($p = 0.39$, ANCOVA test for covariance of the number of strong synapses with wavefront size). When the input firing rate was scaled with the spot size, such that the firing rate was

doubled when the spot size was halved, the RF refined to the same size, regardless of spot size (Fig. 5.6I, grey; $p = 0.14$, t -test for the dependence of the number of strong synapses on wave speed with scaled input firing rates). When the input firing rate was scaled, the final RF sizes under wave activity patterns was the same as for spot activity patterns ($p = 1$, two sample t -test). An example of RF development when the input rate was scaled is provided in Fig. 5.6H, for comparison with Fig. 5.6G.

In summary, we have shown that, in contrast to more traditional implementations of STDP, RF refinement with the STDP_{sub} model is much more sensitive to the total input activity at a given time than it is to the bandwidth of spatial frequencies imposed on the connectivity pattern, both of which may be modulated by the wave speed and burst duration. To this effect, we demonstrated that, with the STDP_{sub} model, the distinct spatiotemporal correlations induced by travelling waves are not necessary to instruct refinement of RFs, and that spatially correlated but stationary spots of activity are sufficient. In the next section, we build upon these results to address the dependence of RF refinement on retinal wave propagation.

Disrupting retinal wave propagation has contrasting effects on RF refinement under two different models of BTDP

A number of experimental attempts to probe the role of retinal waves in the refinement of retinotopy and individual RFs have utilised mouse models in which synaptic transmission, via nicotinic acetylcholine receptors that contain the $\beta 2$ -subunit ($\beta 2$ -nAChRs), is impaired by genetic knockout or pharmacological inhibition of those receptors (Feller et al., 1996; Bansal et al., 2002; Stafford et al., 2009). These manipulations prevent the characteristic propagation of stage 2 retinal waves that depend on $\beta 2$ -nAChR transmission. Rather, these perturbed waves exhibit large wavefronts that travel faster than wild-type waves (Stafford et al., 2009). After such perturba-

tions, the precision with which axons from retinal ganglion cells (RGCs) innervate the SC is dramatically reduced (McLaughlin et al., 2003a; Chandrasekaran et al., 2005; Pfeifferberger et al., 2006; Cang et al., 2008; Huberman et al., 2008a; Dhande et al., 2011). However, in $\beta 2$ -nAChR knockout ($\beta 2^{-/-}$) mice, transgenic expression of $\beta 2$ -nAChRs in RGCs produces small waves that propagate only short distances (Xu et al., 2011). In these mice, the termination of RGC axons in the SC and LGN matches the precision of WT mice in monocular areas but fails to refine in binocularly innervated areas of these nuclei. In this section, we seek additional insight into these experimental results, and the implications they have for the types of plasticity rule that may be at play, by modifying the propagation of retinal waves in the Feller et al. (1997) model and examining the effect this has on RF development in our simulation framework.

Stage 2 retinal waves cast long lasting refractory domains in the starburst amacrine cell (SAC) layer through which future waves cannot propagate. Therefore, in order to manipulate wave propagation in the Feller et al., 1997 retinal wave model, we varied the mean refractory period of starburst amacrine cells (SACs), τ_{ref} (see General Methods), between simulations in the range 40–120 s (SD 20 s), which lies within the broad range of inter-wave intervals and SAC refractory current durations found experimentally (Feller et al., 1997; Zheng et al., 2006; Ford et al., 2012). This had a clear impact on the propagation of waves, examples of which are provided in Fig. 5.7A and B. Waves often swept across large areas of the input layer when τ_{ref} was short (Fig. 5.7A, $\tau_{\text{ref}} = 40$ s) and were much more confined when τ_{ref} was long (Fig. 5.7B, $\tau_{\text{ref}} = 120$ s).

We ran two sets of simulations, one using the traditional implementation of the symmetric STDP rule with a long time constant ($\tau_+ = 0.3$ s), as in the beginning of this chapter, and the other using the STDP_{sub} model ($\tau_+ = 50$ ms, $C_- = 0.003$ /spike).

These two rules had highly contrasting effects on RF development as τ_{ref} , and thus the freedom for wave propagation, was varied between simulations. Examples of final RFs in simulations that used the symmetric STDP rule and the STDP_{sub} model are provided in Fig. 5.7C and D, respectively. With the symmetric STDP rule, simulations in which $\tau_{\text{ref}} = 40$ s exactly matched those used at the beginning of this chapter, producing large sweeping waves that drove the refinement of the RF to a smaller size (Fig. 5.7C, left). However, in simulations that used $\tau_{\text{ref}} > 40$ s, RFs split into subfields, with more subfields appearing as τ_{ref} was made longer, until the RF occupied almost the entire area within the harbour function (Fig. 5.7C, second left to right). We defined the RF area as the total area on the input layer providing synaptic input with strength > 0.5 . In Fig. 5.7E, we summarise the development of the RF area over the course of the simulation for each value of τ_{ref} , averaged over three repeated trials that used different seeds for the random number generators. Fig. 5.7 shows the clear trend for the development of larger RF areas with longer τ_{ref} values, and also shows that the final RF areas were stable, with the exception perhaps of RFs in simulations with $\tau_{\text{ref}} = 120$ s, for which RF growth appeared to be restricted by the harbour. In direct contrast, when the influence of the STDP_{sub} model, RFs maintained a single subfield and refined to smaller areas in simulations that used longer values of τ_{ref} (Fig. 5.7C). The development of the RF area with the STDP_{sub} model, averaged over the three repeated trials, is summarised in Fig. 5.7F.

To elucidate why reduced wave propagation has opposing effects on RF development for the two plasticity rules tested in our simulations, we used the space-time cross-correlogram (sCCG) to measure differences in the spatiotemporal structure of correlations induced by simulated retinal waves with different τ_{ref} values. Waves were isolated from a 5000 s segment of the retinal wave simulation, for every repeated trial and for each τ_{ref} value. Thus, waves isolated from a total duration of 15000 s were analysed for each τ_{ref} . This amounted to 1745 waves when $\tau_{\text{ref}} = 40$ s, 1476 waves

when $\tau_{\text{ref}} = 60$ s, 1151 waves when $\tau_{\text{ref}} = 80$ s, 704 waves when $\tau_{\text{ref}} = 100$ s and 309 waves when $\tau_{\text{ref}} = 120$ s. The sCCG maps the mean correlation strength between the spiking activity of all input neuron pairs as a function of their spatial separation, Δr , and the temporal lag between their spikes, Δt (see General Methods). For travelling waves, the sCCG produces a ridge of strong correlations, whereby the Δr and Δt coordinates along the crest of the ridge have an approximately linear relationship. Examples of the sCCG are provided in Fig. 5.8A for several simulated waves, when $\tau_{\text{ref}} = 40$ s (left column) and $\tau_{\text{ref}} = 120$ s (right column). A wave was classed as ‘travelling’ if a linear regression, shown as a solid purple line in Fig. 5.8, could be well fit ($R^2 \geq 0.9$) to the ridge of its respective sCCG (see General Methods). All other waves were classed as ‘non-travelling’. Regardless of τ_{ref} , all simulations contained waves that were classed as travelling (upper three sCCGs in each column of Fig. 5.8) and non-travelling (lower two sCCGs in each column of Fig. 5.8).

The propagation of a travelling wave cast strong correlations over relatively large distances. These distances can be measured using a proxy, Δr_{max} , which corresponds to the maximum value of Δr to which the linear regression could be fit (indicated by dashed purple lines in Fig. 5.8A). It is important to note that, for the following analysis, Δr_{max} provides a more meaningful description of wave propagation than does, for example, the total path length of the wave, or the total retinal area covered by the wave, which have been typical measures of wave propagation in previous studies, as we are ultimately interested in how wave induced correlations interact with the plasticity rules to instruct the development of the connectivity pattern. In Fig. 5.8B, we plot the cumulative distribution for Δr_{max} from simulations using each τ_{ref} value. Δr_{max} diminished considerably with increasing τ_{ref} ($p < 0.001$, Kruskal-Wallis test for all τ_{ref} , $p < 0.001$ for Δr_{max} distributions from consecutive τ_{ref} values, one-tailed Mann-Whitney U-tests), confirming our initial visual inspection of the simulated retinal waves.

In order to address the influence of the STDP_{sub} model on RF development, we examined the effect of reduced wave propagation on the total input activity per unit time. For all isolated waves determined as both travelling and non-travelling by the sCCG, we measured the average number of spikes per millisecond (the temporal resolution of the simulation) per wave, n_{spk} , and show the cumulative distribution of n_{spk} for each τ_{ref} in Fig. 5.8C. Just as for Δr_{max} , n_{spk} decreased considerably with increasing τ_{ref} ($p < 0.001$, Kruskal-Wallis test for all τ_{ref} , $p < 0.001$ for n_{spk} distributions from consecutive τ_{ref} values, one-tailed Mann-Whitney U-tests). Indeed, for individual waves, Δr_{max} was correlated with n_{spk} across all τ_{ref} values (Fig. 5.8D, $R^2 = 0.4$). With reduced wave propagation came a reduction in the percentage of waves that were classed as travelling (Fig. 5.8E). Non-travelling waves were substantially smaller than travelling waves in simulations with a given τ_{ref} (Fig. 5.8F; $p < 0.01$, one-tailed Mann-Whitney U-test for the equality of medians in n_{spk} distributions between travelling and non-travelling waves, for each τ_{ref} value), further exacerbating the reduction in n_{spk} with reduced wave propagation. These analyses and our results from the 1D simulations above strongly imply that a reduction in the total input activity per unit time, which follows a reduction in wave propagation, gives rise to greater RF refinement under the STDP_{sub} rule with increasing τ_{ref} in our retinal wave simulations.

With the symmetric STDP rule, reduced wave propagation resulted in the RF splitting into subfields and larger RF areas. In Chapter 4, we demonstrated that RF splitting could be a result of slow wave speeds imposing higher spatial frequencies on the connectivity pattern. We therefore tested whether wave speed decreased with increasing τ_{ref} . For every wave that was classed as travelling, wave speeds were computed from the gradient of the linear regression that was fit to the sCCG ridge (see General Methods). In Fig. 5.8G, we show the cumulative distributions of

wave speeds for each τ_{ref} . Very little variation in wave speed with τ_{ref} was observed, although wave speeds when $\tau_{\text{ref}} = 40$ s were significantly faster than wave speeds for longer τ_{ref} ($p < 0.001$, Kruskal-Wallis test for all τ_{ref} ; $p = 0.71$, Kruskal-Wallis test for $\tau_{\text{ref}} = 60$ – 120 s). Nevertheless the median wave speed was only a little faster (~ 0.017 mm/s) when $\tau_{\text{ref}} = 40$ s compared to wave speeds for other τ_{ref} values. Thus, the trend in RF development with increasing τ_{ref} for the symmetric STDP rule must be a consequence of something other than wave speed. We also demonstrated above that the total input activity at a given time has little effect on the development of the connectivity pattern with this traditional STDP rule.

We therefore hypothesised that the reduction in wave propagation with increasing τ_{ref} violated one of the assumptions made in Chapter 3 that was necessary to derive the analytical description of synaptic strength dynamics in Eq. 3.6. This assumption stated that the duration of each wave must be sufficiently longer than the time scale of the STDP rule. An equivalent interpretation of this assumption, for a wave travelling with a constant speed, is that the wave must propagate a distance greater than $1/k^*$, where k^* is the characteristic spatial frequency of $\kappa(x)$. In Fig. 5.8H, we demonstrate how this assumption is increasingly violated by waves that were generated in simulations using longer values of τ_{ref} . The grey curve shows an estimate of $\kappa(x)$ that was computed using $\tau_+ = 0.3$ s and $v = 0.21$ mm/s, which was the median wave speed in simulations that used $\tau_{\text{ref}} = 40$ s. Against this curve, we plot the median measures for Δr_{max} from simulations using different τ_{ref} values (dark green, green and bright green for τ_{ref} values of 40 s, 80 s and 120 s, respectively). Because synaptic strength dynamics can be approximated by the inner product of the cross-correlogram and the STDP rule (Kempster et al., 1999), these measures of Δr_{max} indicate how $\kappa(x)$ is sampled by waves that propagate over shorter or longer distances. Thus, as τ_{ref} is increased, Δr_{max} decreases and travelling waves sample an ever narrower window of $\kappa(x)$. When $\tau_{\text{ref}} = 40$ s, more than half of all travelling

waves sample a window that is wider than the full width of $\kappa(x)$, whereas, when $\tau_{\text{ref}} = 80\text{ s}$, only half of all travelling waves sample the full width of $\kappa(x)$, and when $\tau_{\text{ref}} = 120\text{ s}$, more than half of all travelling waves sample a window that is narrower than the full width of $\kappa(x)$. Thus, as τ_{ref} is increased, a larger percentage of travelling waves violate the assumption that is necessary for the interaction of travelling waves with the STDP rule to be described by Eq. 3.6 in Chapter 3. This is further exacerbated by the increase in the percentage of non-travelling waves with τ_{ref} .

We therefore require a new description for the dynamics of synaptic strengths under the influence of waves for which propagation is reduced or abolished. Such waves are short lived, and all input neurons that are recruited by such waves are recruited almost simultaneously, which is apparent in the sCCGs for non-travelling waves in Fig. 5.8A, in which spatiotemporal correlations are concentrated around the origin. Because the symmetric STDP rule specifies synaptic strengthening for short Δt , these waves should always lead to the strengthening of corresponding synapses, so long as the output neuron is currently active. Thus, the trend in the final RF properties with increasing τ_{ref} can be explained by the corresponding trend in the synaptic strength dynamics. When $\tau_{\text{ref}} = 40\text{ s}$, the vast majority of waves propagate over relatively large distances and give rise to pattern formation as described by Eq. 3.6 in Chapter 3 and RF refinement as demonstrated in Chapter 4 and in Fig. 5.2. When $\tau_{\text{ref}} = 60\text{--}80\text{ s}$, the increased percentage of waves that travel only short distances, or that are non-travelling, heightens the probability that additional islands of synapses can be strengthened because they can drive the output to spike within the window for synaptic strengthening. However, in this parameter regime, there may remain a sufficient number of waves that travel over long distances. Consequently, the passage of larger waves over a nascent subfield reinforces the subfield and weakens the connectivity pattern in between subfields. Thus, a RF consisting of multiple subfields becomes a stable solution for the synaptic strength dynam-




Figure 5.7: **The disruption of wave propagation affects RF development in a manner that depends on the mechanisms underlying BTDP.** **A.** Snapshots of typical waves that were isolated in the retinal wave simulation when $\tau_{\text{ref}} = 40\text{ s}$ (**A**, 1s between consecutive snapshots) and $\tau_{\text{ref}} = 120\text{ s}$ (**B**, 0.2s between consecutive snapshots). Input neurons that fired at least one spike in a window of 0.1s are black. **C-D.** Example RFs that developed when using either the symmetric STDP rule (**C**) or the STDP_{sub} model (**D**) in simulations using different τ_{ref} values. Solid green circle: harbour function boundary. Dashed green circle: initial RF boundary. **E-F.** Summary for the development of the RF area over the course of the simulation. Dark to light green: τ_{ref} values from 40s to 120s. Each curve is an average over 3 repeated trials. Lower dashed line: initial RF area. Upper dashed line: area of the harbour. **E.** Development of RF areas under the symmetric STDP rule. **F.** Development of RF areas with the STDP_{sub} model.




Figure 5.8: The spatial and temporal properties of retinal waves reveals how disrupted wave propagation interacts with different implementations of BTDP. **A.** Examples of the sCCG for five isolated waves in retinal wave simulations with $\tau_{\text{ref}} = 40$ s (left column) and $\tau_{\text{ref}} = 120$ s (right column). Solid purple lines: linear regressions fit to the crest in the ridge of strong correlations. Regressions are shown only in the region of the ridge to which they were fit. Dashed purple lines denote $\Delta r_{\max}$. Temporal and spatial scales are provided by the scale bars. $\Delta t = 0$ and $\Delta r = 0$ lie at the bottom left corner. **B.** Cumulative distributions for wave propagation distances, as measured by $\Delta r_{\max}$. Dark to light green: τ_{ref} values from 40 s to 120 s. **C.** Cumulative distributions for total input activity per unit time per wave, as measured by n_{spk} , for all isolated waves. Dark to light green: τ_{ref} values from 40 s to 120 s. **D.** Correlation between $\Delta r_{\max}$ and n_{spk} measurements for all waves classed as travelling. The τ_{ref} corresponding to the simulation from which a wave was analysed is colour coded as in **B** and **C**. Red line: linear regression. **E.** The percentage of waves that were classed as travelling as a function of τ_{ref} . **F.** Non-travelling waves always produce less input activity per unit time per wave than do travelling waves. **G.** Cumulative distributions for wave speeds, as measured by the gradient in linear regressions of the ridge in sCCGs. Dark to light green: τ_{ref} values from 40 s to 120 s. **H.** Grey: estimate of $\kappa(x)$ for the symmetric STDP rule, using $v = 0.21$ mm/s, $\tau_+ = 0.3$ s and a burst duration of 1 s. The median measurements of $\Delta r_{\max}$ are marked against $\kappa(x)$ for simulations in which $\tau_{\text{ref}} = 40$ s (dark green), $\tau_{\text{ref}} = 80$ s (medium green) and $\tau_{\text{ref}} = 120$ s (bright green). These indicate the maximum spatial windows within which $\kappa(x)$ is sampled by 50% of the waves that are classed as travelling.

ics. Increasing τ_{ref} to 100–120 s further increases the percentage of non-travelling and spatially restricted travelling waves. In this regime, the emergence of new subfields, due to the influence of small waves, outweighs the ability of waves that travel larger distances to impose synaptic weakening, resulting in RFs that engulf the area available to them, i.e. almost all synapses within the arbour obtain the maximum strength.

Discussion

The role for retinal waves in the refinement of the retinocollicular circuit is far from clear. The induction of localised correlations by retinal waves, and the disruption of retinotopy and RF refinement in the SC when retinal waves are perturbed, has led to the hypothesis that wave-induced correlations instruct refinement through a Hebbian mechanism of plasticity. At the retinogeniculate synapse, Hebbian plasticity takes the form of a BTDP rule (Butts et al., 2007), which operates over seconds rather than tens of milliseconds, as predicted by Butts & Rokhsar (2001). As yet, the underlying mechanism by which BTDP rules arise, and whether this rule also operates at retinocollicular synapses, is unknown. Here, we have examined two possible types of STDP rule that also yield the BTDP rule. First, we showed that a phenomenological STDP rule (Song et al., 2000), which operates over time scales lasting hundreds of milliseconds, can engage the space-time inseparable correlations induced by retinal waves to refine RFs by imposing upon them a specific spatial frequency. In contrast, the STDP_{sub} model, which we used as a proxy for a more mechanistic, $[\text{Ca}^{2+}]$ dependent model of STDP (Graupner & Brunel, 2012), is sensitive only to spatial correlations and the summed activity arriving at the postsynaptic neuron, such that RF refinement is independent of wave propagation. Thus, these two different plasticity rules assign importance to different features of retinal waves. Under the traditional implementation of the STDP rule, the wave speed, STDP

time scale and burst duration are influential features, but the input firing rate is not. Moreover, there is a minimum distance over which waves must propagate for the STDP rule to instruct pattern formation in the synaptic strengths. Under the STDP_{sub} rule, the wave speed, burst duration and input firing rates are influential features, but the STDP time scale is not, at least within a realistic range. Yet there is no restriction on the distance travelled by the wave for this rule to instruct development of the connectivity pattern.

Plausibility of STDP rules with long time scales

A number of lines of evidence support the possibility of plasticity rules operating over longer time scales during early postnatal stages of development. STDP in many cases depends on the influx of Ca^{2+} through NMDARs (Feldman, 2012; Shipton & Paulsen, 2014), which, owing to their dominant composition of NR2B subunits during stage 2 retinal waves, have long decay times (Shi et al., 1997; Liu et al., 2004). The influence of NMDAR subunit composition on STDP time scales has recently been tested, providing evidence that NR2A and NR2B subunits support STDP time scales of ~ 50 ms and 150 ms, respectively (Guo et al., 2012). The precise mechanisms underlying this form of plasticity, and how it may play out for longer and more complex spike sequences as are found during retinal waves, are yet to be addressed. Nevertheless, these findings demonstrate the plausibility of STDP rules that operate over time scales considerably longer than have previously been reported.

STDP decay times of ~ 150 ms are, however, still too short to impose the correct spatial frequency on developing RFs in SC, as predicted by our analytical results. How might longer time scales be obtained? The Graupner & Brunel (2012) $[\text{Ca}^{2+}]$ model, with a $[\text{Ca}^{2+}]$ decay time of 20 ms, gave rise to a wide window for LTP, due to the overlapping and therefore correlated input and output bursts. However, one limitation of the $[\text{Ca}^{2+}]$ model is that it is sensitive to correlated activity only insofar

as the resulting summed $[Ca^{2+}]$ transients can exceed the thresholds for LTP and LTD. It is therefore not strictly a correlation based plasticity rule, i.e. it does not solely depend on the product of input and output activity. As such, the summed $[Ca^{2+}]$ transients were very sensitive to small fluctuations in the input and output firing rates, such that the effective BTDP produced by the $[Ca^{2+}]$ model was laden with noise. It is therefore possible that variable firing rates could easily occlude the influence of spatiotemporal correlations on plasticity. A second limitation is that the summed $[Ca^{2+}]$ acts as the only coincidence detector, whereas a number of studies have demonstrated that separate coincidence detectors are needed for LTP and LTD (Karmarkar & Buonomano, 2002; Sjöström et al., 2003; Bender et al., 2006). A consequence of only one coincidence detector in the $[Ca^{2+}]$ model, with $[Ca^{2+}]$ decay times of 20 ms, is that the model is unable to produce LTD over time scales that are > 200 ms yet finite, which would be required for the band pass property necessary for periodic patterning. Therefore, the only viable way for the $[Ca^{2+}]$ model to produce LTD over time scales longer than the $[Ca^{2+}]$ decay time, and still yield LTP for overlapping bursts, is for isolated input or output spikes to produce LTD. Implementing this realisation of the $[Ca^{2+}]$ model in the $STDP_{sub}$ model demonstrated that LTD by unpaired output spikes could not yield the required bandpass feature. Thus, periodic patterning could not be achieved, and the travelling nature of the wave served only to influence the spatial scale of correlated input activity. Two coincidence detectors may therefore provide all the components necessary for a viable STDP or BTDP rule that can induce periodic patterning. This resonates with the phenomenological STDP rule, which can be implemented using two variables that separately track the input and output activity to determine the magnitude of LTP and LTD, respectively (Song et al., 2000). The role of endocannabinoid (eCB) signalling is a well documented candidate for a second coincidence detector, whereby eCB signalling gates LTD only (Sjöström et al., 2003; reviewed in Heifets & Castillo (2009) and Feldman (2012)). This process requires postsynaptic activity dependent

release of eCB at the synapse, and coincident opening of presynaptic NMDARs by the presynaptic release of glutamate. Importantly, eCB signals decay slowly and, after bursts of postsynaptic spikes, can lead to LTD windows lasting hundreds of milliseconds (Sjöström et al., 2003). The coincidence detector for LTP can then depend on the time scale of $[Ca^{2+}]$ transients, which is determined by the duration of the pre- and postsynaptic bursts.

Features of retinal waves that may regulate refinement

It is very possible that retinofugal refinement is instructed by a mechanism that is not sensitive to the specific spatiotemporal correlations induced by retinal waves, but to some component feature of those correlations, or to some other feature altogether. For example, retinal waves induce spatially localised correlations which may alone be important, such that their propagation may simply be an epiphenomenon. Spatially localised correlations may act through Hebbian mechanisms of plasticity that depend on the firing rate and synaptic competition, such as with the STDP_{sub} model. A role for spatial correlations in refinement have been modelled extensively, and are reviewed in Simpson et al. (2009). In these models, spatially correlated input activity patterns preserve information about retinotopy, reinforce adjacent synapses and induce competition between distant inputs that are weakly correlated, thus providing a means of refinement. We have demonstrated these same principals using the STDP_{sub} model in Figs. 5.6 and 5.7. In $\beta 2^{-/-}$ mice, which exhibit fast and broad stage 2 retinal waves (Stafford et al., 2009), RGC axons form exuberant arbours in the SC that terminate over diffuse areas in comparison with wild-type (WT) mice, thus disrupting the retinotopic map and leading to abnormally large RFs (McLaughlin et al., 2003a; Chandrasekaran et al., 2005; Pfeifferberger et al., 2006; Huberman et al., 2008a). These studies suggested an instructive role for spatial correlations in stage 2 retinal waves, although they did not confirm whether spatial correlations alone or space-time inseparable correlations are important for

refinement. However, this interpretation is confounded by the lack of $\beta 2$ subunit containing nAChRs throughout the brain, not just in the retina. To address this issue, Xu et al. (2011) expressed a transgene for $\beta 2$ subunit nAChRs in the retinae of $\beta 2^{-/-}$ mice. In these $\beta 2^{-/-}$ transgenic ($\beta 2^{-/-}$ (TG)) mice, the normal refinement of RGC arbours in the SC was rescued. However, stage 2 retinal waves recorded in these mice propagated over much smaller distances than in WT mice. As we showed in Fig. 5.8, retinal waves that did not travel over large enough distances failed to impose periodic patterning on the feedforward network and refine RFs. Moreover, many of these waves failed to exhibit distinct space-time inseparable correlations at all. Whether waves in $\beta 2^{-/-}$ (TG) mice are too small, at least by the standards described using our sCCG, to instruct refinement in this manner, will require further analysis. The outcome of such an analysis would help to establish whether wave propagation is required for a Hebbian mechanism of refinement, or whether spatially local correlations alone are sufficient.

The temporal properties of retinal wave generation may also carry information. The dynamic processes that generate retinal waves mean that a single RGC participates in a wave approximately once every 1–3 minutes (Feller et al., 1997; Stafford et al., 2009), resulting in oscillations of intracellular cyclic AMP (Dunn et al., 2006) that may be required for RGC sensitivity to the repulsive signals between the ephrin-A and EphA family of molecular guidance cues (Nicol et al., 2007). The temporal properties of retinal waves also help to reduce correlated activity between the eyes, triggering debate surrounding a role for retinal waves in segregating eye specific inputs to SC and LGN (Feller, 2009; Chalupa, 2009). Whether the temporal dynamics of cyclic AMP influences RGC axon targeting *in vivo* is yet to be determined.

Experimental predictions and future directions

Based on our results, we identified three properties of retinal waves that are pre-

dicted by our models to have a quantifiable effect on retinocollicular development in a manner that differentiates between the two plasticity rules. For both rules, faster wave speeds are expected to yield larger RFs. However, waves that propagate over smaller areas are predicted to produce larger RFs under the phenomenological STDP rule, and to produce smaller RFs under the STDP_{sub} rule. Finally, a STDP_{sub} rule should produce larger RFs for retinal waves that exhibit greater RGC firing rates, whereas the firing rate should have a negligible effect under the phenomenological STDP rule.

Addressing our predictions quantitatively will require a high degree of control over the pattern of retinal waves *in vivo*. Many studies that have attempted to manipulate retinal wave activity have utilised the $\beta 2^{-/-}$ mouse model, but it has proven difficult to formulate precise interpretations about their results, in part due to the variability of wave characteristics in different experimental preparations, but also because of potential side effects on the expression of genes associated with cell adhesion and calcium signalling with this genetic model (Sun et al., 2008a; Rubin et al., 2011; reviewed in Cang & Feldheim, 2013). Recent attempts to control retinal activity during early development have utilised mice that express channelrhodopsin in a subset of RGCs, which were stimulated by illumination with blue LEDs (Zhang et al., 2012). Repetitive pan-retinal LED illumination was provided to the retinae every 5 s during waking hours for two days. As a result of stimulation, the disruption of retinotopy that is common in $\beta 2^{-/-}$ mice was partially mitigated. This is a surprising result, as it conflicts with the hypothesis that localised spatial correlations induced by retinal waves instruct refinement. However, unreliable post-synaptic responses in the SC after pan-retinal illumination cast doubt on whether RGCs throughout the retina were reliably driven to spike. A notable difficulty with an experiment like this is the lack of transparency in the eyeball optics at this age (Jester, 2008), making controlled stimulation of retinal waves with this method a

considerable challenge.

A meaningful test of our experimental predictions will also require greater understanding of synaptic plasticity at the retinocollicular synapse, and its dependence on realistic pairings of wave-induced pre- and postsynaptic bursts. We propose that, if the space-time inseparable correlations of travelling waves are required to instruct retinocollicular refinement, then a STDP rule at this synapse should have the properties of a bandpass filter and should be sensitive to the coincidence of pre- and postsynaptic spikes over time scales lasting hundreds of milliseconds to seconds. To help elucidate whether bandpass-like STDP rules exist at retinocollicular synapses during stage 2 retinal waves, a standard experimental protocol in which the timing dependence of STDP is deduced could be conducted. This would entail pairing realistic bursts in connected RGC and SC neurons over a sufficiently wide range of temporal offsets, and for different burst durations. However, if synaptic plasticity lacks the bandpass property, as would be expected if governed by the Ca^{2+} dynamics modelled in Graupner & Brunel (2012), then the travelling nature of retinal waves are an epiphenomenon of some other, important feature of activity, as yet to be determined.

6. Results

The influence of delayed inputs on pattern formation and its application the development of simple cells in primary visual cortex

Background

The primary visual cortex (V1) is the first stage of processing along the visual pathway for a number of visual features (Carandini et al., 2005). Selectivity for these features continues to refine, and in some cases is first established, during early postnatal life, and is dependent on visual experience (Espinosa & Stryker, 2012), making V1 highly amenable to the theoretical and experimental investigation of RF refinement. Two such features are extracted by receptive fields (RFs) belonging to a class of neuron called simple cells, which will become the focus of this chapter. These features include: 1) selectivity for the orientation of luminance contours, as a result of spatially offset subfields that respond to either increases or to decreases in

luminance, respectively called ‘ON’ and ‘OFF’ responses; 2) selectivity for motion in a particular direction, which is perpendicular to the preferred orientation (Hubel & Wiesel, 1959; Hubel & Wiesel, 1962). The similarity in responses between ON and OFF subfields in simple cells and ON and OFF response properties of neurons in the LGN prompted Hubel and Wiesel to hypothesise that simple cells acquire the spatial layout of their RF through the summation of feedforward inputs from the LGN, which was supported by later studies (Reid & Alonso, 1995; Usrey et al., 2000; De Valois et al., 2000; Alonso et al., 2001). Direction selectivity, on the other hand, requires inputs that are offset in both space and time (Hubel & Wiesel, 1962; Barlow & Levick, 1965). However, the source of these spatial and temporal offsets in V1 is still uncertain (Baker & Bair, 2012).

While orientation selectivity is already established at the time of eye opening (Chapman & Stryker, 1993; Krug et al., 2001; White et al., 2001; Li et al., 2006), the maturation of direction selectivity in simple cells requires exposure to motion in patterned visual environments; prolonged exposure to strobe lighting environments, for example, prevents the development of direction selective (DS) RFs (Humphrey & Saul, 1998). The importance of exposure to visual motion has been emphasised in a number of studies in which animals were reared in environments biased for motion in one direction. Unlike normally reared animals, for which DS cells in V1 express selectivity over all directions, DS cells in animals that were reared in biased motion environments expressed a corresponding bias in their preferred directions (Cynader et al., 1975; van Hooser et al., 2012). These findings suggest that the development of the spatial and temporal organisation of inputs to DS cells in V1 is influenced by the spatial and temporal structure of the visual environment. How this process unfolds, and which features of the environment are most informative for development, is far from clear. One candidate model of direction selectivity posits that both spatial and temporal offsets are seeded by afferents from the LGN (Saul

& Humphrey, 1990; de Valois & Cottaris, 1998; De Valois et al., 2000; Saul et al., 2005). From early ages, there exist a range of response latencies in LGN neurons (Daniels et al., 1978; Blakemore & Vital-Durand, 1986a; Tavazoie & Reid, 2000), thus providing the necessary temporal delays between inputs, but how these inputs become organised in space during the development of direction selectivity in V1 is unknown.

Models of linear Hebbian plasticity have made useful contributions to this question. Miller (1994) proposed that negative correlations between overlapping ON and OFF RFs in the LGN, and positive correlations between spatially offset ON and OFF RFs, can drive the spatial segregation of ON and OFF inputs to a simple cell via a linear Hebbian plasticity rule. In addition, the same plasticity rule can strengthen spatially offset non-lagged and lagged inputs that are exposed to visual stimuli with motion biased for one direction (Wimbauer et al., 1997b). The spatial offset in this case is proportional to the distance travelled by the stimulus during the delay between non-lagged and lagged inputs. By this process, linear models readily develop the RF profile of DS simple cells when exposed to moving visual stimuli (Wimbauer et al., 1997a). However, without biased motion, such linear models learn only the mean temporal pattern of correlations and fail to develop selectivity for any one direction. A simple solution is therefore to add a non-linearity that breaks the symmetry of synaptic strength changes in response to motion in different directions (Feidler et al., 1997). Evidence against these models came from Ohshiro & Weliky (2006), who showed that, at an age when direction selectivity is still maturing, correlations between the activity of ON and OFF cells in the LGN do not exhibit the spatial structure proposed by Miller (1994). Moreover, in all of these models, the spatial structure of the simple cell RF is determined only by spatial correlations between the inputs, and not by temporal correlations beyond coincident activity. As a consequence, the spatial offset between non-lagged and lagged inputs that produces

strong DS responses in the simple cell can only be achieved with motion stimuli that have a restricted range of velocities (Blais et al., 2000), which is not characteristic of natural visual environments (Dong & Atick, 1995).

Attempts to develop more biophysically realistic models of synaptic plasticity may help to address the shortcomings of previous, rate based Hebbian models. Spike-timing dependent plasticity (STDP) is a good candidate plasticity rule, due to its sensitivity to the temporally correlated activity between inputs. Temporally asymmetric manifestations of STDP shift the spatial profile of synaptic strengths in models that are driven by stimuli with biased motion (Rao & Sejnowski, 2003; Gütig et al., 2003; Wensch et al., 2005; Honda et al., 2011). Coupled with spatially offset lagged populations, STDP can also build DS simple cell RF profiles (Buchs & Senn, 2002). However, the simple cells in Buchs & Senn (2002) reliably develop DS responses only if some initial spatial organisation of LGN inputs is provided from the outset. Moreover, different response delays between inputs in the Buchs & Senn model are achieved through synapses that undergo short term depression (Chance et al., 1998), a method that is only effective for visual stimuli with periodic luminance modulation.

In this chapter, we examine how our results in previous chapters may apply to the development of simple cells in V1. We extend our models of STDP with traveling waves, which comprise a single input population, to models with multiple input populations that are intrinsically delayed with respect to a moving stimulus, as in the LGN, and that provide synaptic input to a common output neuron. We present analytic solutions to the equations of synaptic strength dynamics in this extended model and support these results in a simulated network of spiking neurons. In particular, we find that periodic connectivity patterns develop for each input population, and that the connectivity pattern from one population develops a phase shift,

with respect to the connectivity pattern of a second population, that is independent of the stimulus speed and determined only by properties of the circuit, including the STDP rule, stimulus response functions and the relative delay between the two populations. We interpret our results in the context of a developing geniculocortical circuit consisting of ON, OFF, non-lagged and lagged inputs that is exposed to moving luminance contours, and demonstrate the likeness of mature output neurons in our model to DS simple cells in V1. In particular, we find that STDP rules, such as those found in the developing V1 area (Feldman, 2000), may be tuned to seed simple cells in a manner that jointly optimises direction selectivity and response amplitude. Furthermore, we predict that a range of phase shifts between ON and OFF inputs should be observed in geniculorecipient cells in V1. We discuss these results as a foundation on which to build further models of simple cell development that may help to overcome the shortcomings of previously published models.

Methods

Spiking neuron simulation with multiple input populations

Details concerning the architecture and the general dynamics of the simulation can be found in General Methods. Here, we describe details specific to this chapter and extensions to the simulation to incorporate multiple input populations.

Each input population was constructed in the same manner as for the single population simulations in previous chapters. Thus, each population had its own set of excitatory synapses that converged onto a single and common output neuron. There were no synaptic connections between input neurons in different populations. Excitatory postsynaptic potentials (EPSPs) elicited by spikes from each input population were summed linearly by the output neuron. Individual populations in the

input layer consisted of 500 input neurons arranged along a one dimensional lattice with a lattice spacing of 0.5° . Although a single mammalian eye does not see beyond 180° of visual angle, we use this comparatively large input layer so as to accurately measure the spatial frequency of periodic connectivity patterns that develop. We considered wavefronts in the input layer to be a response to a training stimulus, in this case, a moving, high contrast spot. Training stimuli moved at a speed of $50^\circ/\text{s}$ (as in Li et al., 2008), recruiting input neurons as it passed them. However, as we show in the Results, the development of direction selectivity in the output is independent of the training stimulus speed. When recruited, an input neuron underwent a burst of spikes for 100 ms with a firing rate of 50 Hz, indicative of the slow light responses in young animals (Blakemore & Vital-Durand, 1986a; Cai et al., 1997; Tavazoie & Reid, 2000). Travelling wavefronts in each input population were generated in the same manner, and in all simulations travelled from the left to the right side of the input layer. The only difference between populations was the latency, δt , at which the wavefront was initiated with respect to the stimulus onset. For simplicity, the delay between successively active populations was set by an integer multiple of t_d seconds, whereby the first population responded immediately. Thus, the k^{th} population responded with a latency of $\delta t = kt_d$, for $k = 0, 1, \dots$. The interval between successive training stimuli was sufficiently long to allow waves in all populations to have traversed the input layer before the next stimulus arrived.

Calculating phase shifts between connectivity patterns

The phase shift between connectivity patterns, $w_m(x)$ and $w_n(x)$, which correspond to input populations m and n , can be obtained by measuring the shift in spatial coordinates in degrees, $D_{\text{sim}}(w_m, w_n)$, and the dominant spatial frequency of the connectivity patterns in cycles per degree, k_{sim}^* , and taking their product. We use the subscript ‘sim’ to denote variables that were measured from the connectivity patterns that developed in each simulation, as opposed to the predicted value of the

variable that was computed numerically. We measured $D_{\text{sim}}(w_m, w_n)$ by locating the peak in the normalised cross-correlogram (CCG) between w_m and w_n . The cross-correlation, C_{mn} , between connectivity patterns m and n is given by:

$$C_{mn}(\Delta x) = \sum_{x=1}^X w_m(x + \Delta x)w_n(x), \quad (6.1)$$

where X is the size of the input layer. Normalisation is then performed to take into account the extent of overlap between w_m and w_n using a triangle function, $N(\Delta x) = X - |\Delta x|$. The normalised CCG is:

$$CCG(\Delta x) = \frac{C_{mn}(\Delta x)}{N(\Delta x)}. \quad (6.2)$$

The CCG was smoothed using a Gaussian filter, with a SD of 2.5° , before $D_{\text{sim}}(w_m, w_n)$ was extracted.

The spatial frequency used to calculate the phase shift was taken as the mean, dominant spatial frequency over all connectivity patterns in the simulation, $k_{\text{sim}}^* = \sum_m k_{\text{sim},m}^*/M$, where M is the number of input populations (either two or four, depending on the simulation) and $k_{\text{sim},m}^*$ is the centre of a Gaussian curve that was fit to the spatial power spectra of w_m . The measured phase shift, $\Delta\phi_{\text{sim}}$, was therefore $\Delta\phi_{\text{sim}} = D_{\text{sim}}(w_m, w_n)k_{\text{sim}}^*$.

Construction of spatiotemporal receptive fields

To address how direction selectivity in V1 may arise from developing geniculocortical circuits, we designated response characteristics that are typical in the LGN to each input population, in simulations in which there were four populations, and analysed the spatiotemporal RF (STRF) resulting from the combined connectivity patterns with the output neuron. Characteristic response types included ON and OFF, which

respectively respond to local increases and decreases in luminance, and non-lagged (N) and lagged (L), which differed in the latency of the response with respect to the stimulus. STRFs, $G(x, t)$, were constructed by taking the outer product of the spatial connectivity patterns from the simulations, $w_A^B(x)$, with the stimulus response function, $h_A^B(t)$, and summing over the four input types:

$$G(x, t) = \sum_A \sum_B w_A^B(x) h_A^B(t), \quad (6.3)$$

where ‘A’ denotes either ON or OFF and ‘B’ denotes either N or L. The stimulus response describes the synaptic activation, averaged over many stimuli, provided to the output neuron, due to the response of a single input neuron of a given response type:

$$h_{\text{ON}}^{\text{N}}(t) = \alpha(t) * \epsilon(t) \quad (6.4)$$

$$h_{\text{ON}}^{\text{L}}(t) = \alpha(t - t_d) * \epsilon(t) \quad (6.5)$$

$$h_{\text{OFF}}^{\text{N}}(t) = -\alpha(t) * \epsilon(t) \quad (6.6)$$

$$h_{\text{OFF}}^{\text{L}}(t) = -\alpha(t - t_d) * \epsilon(t), \quad (6.7)$$

where $\alpha(t)$ describes the stimulus response of the input firing rate (*boxcar* function with width of 100 ms) and $\epsilon(t)$ describes the EPSP (difference of exponentials; details in General Methods). Both $\alpha(t)$ and $\epsilon(t)$ take the same form here as in the travelling wave simulations, and have a temporal resolution of 1 ms. The STRF therefore describes the change in firing rate of an output neuron in response to a stimulus comprising a change in luminance at position x and time t .

Response of output neurons to test stimuli

To examine the direction selectivity of an output neuron with a given STRF, we computed its linear response to drifting sinusoidal test stimuli. The luminance of

the test stimulus at position x and time t is given by:

$$s(x, t) = \sin(2\pi(xf_x - tf_t)), \quad (6.8)$$

where f_x is the spatial frequency and f_t the temporal frequency of the sinusoidal grating. The firing rate of a neuron with a linear STRF is the convolution in time of the stimulus with the STRF. This equates to multiplying the Fourier transforms of the STRF and stimulus in frequency space, where the Fourier transform of $s(x, t)$ is a delta function, $\delta(f_x, f_t)$. The amplitude of the product is equal to the firing rate modulation amplitude in response to the stimulus, which quantifies the sensitivity of the STRF to the spatiotemporal features of the stimulus. Thus, the response amplitude of the output neuron to drifting sinusoidal gratings over all spatial and temporal frequencies is given by the amplitude spectrum:

$$A(f_x, f_t) = (\text{Re}[\tilde{G}]^2 + \text{Im}[\tilde{G}]^2)^{\frac{1}{2}}, \quad (6.9)$$

where $\tilde{G}$ is the Fourier transform of the STRF.

Note that, in order to construct a linear STRF for the output neuron, we have assumed that the synaptic input also responds linearly to the stimulus. As such, when a stimulus of one polarity, either increasing or decreasing luminance, falls on an input neuron of the opposite polarity, respectively OFF or ON, the synaptic activation of the output neuron is effectively inhibitory. Our approach therefore assumes that the output firing rate can be negative, or that it has a non-zero baseline firing rate. However, the baseline firing rate of input neurons in the simulation is 0 Hz, undergo increases in firing rate in response to the same stimulus type only (ON inputs respond to luminance increases, OFF inputs to luminance decreases), and have zero response to stimuli of the opposite stimulus type. Thus the inputs provide

only excitatory synaptic activation to the output neuron. Thus, the framework behaves linearly but with a firing rate threshold at zero. The output neuron in such a linear-threshold circuit therefore exhibits a non-zero mean firing rate when presented with a drifting grating stimulus. Nevertheless, the firing rate modulates about this mean, as it would do in a linear circuit without a threshold. Because our analysis of direction selectivity centres around the modulation amplitude of the firing rate, and not the absolute firing rate, both linear and linear-threshold frameworks yield nearly equivalent results.

Direction selectivity index

Our primary interest in the STRF is its selectivity for motion of the sinusoidal grating stimulus in a particular direction. This is encoded in the amplitude spectrum of the STRF, in which the quadrant with positive spatial and temporal frequencies (quadrant 1) quantifies the response amplitudes for motion in the rightwards direction (same direction as the training stimulus), and the quadrant with negative spatial and positive temporal frequencies (quadrant 2) quantifies the response amplitudes for leftwards motion (Adelson & Bergen, 1985). For each STRF, we compute a direction selectivity index (DSI) using the following:

$$\text{DSI} = \frac{r_1 - r_2}{r_1 + r_2} \tag{6.10}$$

where r_1 and r_2 are the maximum amplitudes in quadrants 1 and 2 of the amplitude spectrum, respectively.

Artificially constructed STRFs

To examine how the direction selectivity of a STRF depends on the spatial and temporal organisation of its four input types, we built models, which we call artificial

STRFs, of the STRFs that formed in the simulations, allowing us to manipulate the spatial and temporal organisation of ON, OFF, N and L inputs. Artificial STRFs were constructed by summing over four spatiotemporal components for each of the input populations. The spatial structure of each component comprised four cycles of a square wave with a spatial resolution of 30 pixels per cycle. The spatial component was multiplied with the stimulus response functions described in Eqs. 6.4–6.7 to obtain the four input types, ON-N, ON-L, OFF-N and OFF-L, with a temporal resolution of 1 ms. These four components were then summed together, as in Eq. 6.3, with different spatial and temporal offsets to build the artificial STRFs. Specifically, we independently varied the spatial shift between ON and OFF inputs and between N and L inputs through one cycle of the square wave, and the temporal delay was varied between N and L inputs through 0–200 ms. For each configuration of the STRF, we computed two values: the DSI (Eq. 6.10) and the maximum modulation amplitude in the preferred direction:

$$\text{MR} = \max(r_1, r_2). \quad (6.11)$$

Results

Outline of synaptic strength dynamics with multiple, delayed input populations

We build upon the framework developed in Chapter 3 to examine how delays influence the development of periodic connectivity patterns by the interaction of travelling waves with STDP. In this new model, we divide the input layer into subpopulations that comprise neurons that have the same response delay. A schematic for the architecture and dynamics of the model is shown in Fig. 6.1A. We simplify matters initially by considering only two input populations, which we shall denote ‘non-lagged’ (N, red cells in Fig. 6.1A) and ‘lagged’ (L, blue cells in Fig. 6.1A),

whereby wavefronts (denoted by the dark coloured cells) in the L population are delayed by t_d seconds with respect to the N population. We will later consider arbitrary numbers of input populations. This network resembles the geniculocortical circuit, in which relay cells in the LGN can be categorised by their response latencies (Saul & Humphrey, 1990). As such, we can think of wavefronts as visual responses that track moving visual stimuli, as demonstrated in primate retinae by Frechette et al. (2005). Keeping with the visual system analogy, the j^{th} neuron in every population responds to a stimulus at the same location in visual space, and the response properties in each population are equal in all respects bar the latency. We further simplify our analysis by considering waves that travel in only one direction, from left to right, with the intention to later address the influence of biased motion on the development of direction selectivity. In the following, we provide a descriptive account of the developing connectivity pattern from two input populations, which we later support with a formal derivation of the synaptic strength dynamics.

To aid our description, we refer to a schematic, in Figs. 6.1B–D, of the developing N ($w_N(x)$, dark red) and L ($w_L(x)$, dark blue) connectivity patterns as a single wavefront traverses the N and L input populations (red and blue circles, respectively). For this description, synapses are modified by a symmetric STDP rule (Eq. 2.4 in General Methods), which does not lead to phase shifts in the connectivity pattern under the influence of unidirectional waves, as does the asymmetric rule. This simplifies our description, but we will show later that the choice of STDP rule does not influence the principal results in this chapter. We first consider the scenario in which the inter-population delay is $t_d = 0$. In this case, the wavefront is at the same location in both populations at any given time: the input layer effectively has only a single population and the synaptic strength dynamics are therefore equivalent to Eq. 3.6 in Chapter 3. Consequently, both input populations will develop the same connectivity pattern, which is periodic in space, with the output neuron. These

connectivity patterns are depicted by the square waves in Fig. 6.1B–C, indicating periodic bands of strong and weak synapses. We use this scenario as a starting condition from which to describe the development of the connectivity patterns after setting the delay in the L population to $t_d > 0$.

To illustrate the pattern of synaptic strength changes as a wave traverses the input layer when $t_d > 0$, we depict a snapshot of the model dynamics when the wavefront in the N population (dark red cell) is at the centre of a peak in the N connectivity pattern, shown in Fig. 6.1B. At this instant, the output neuron will be driven by activity in the N population only because, at the location of the L wavefront, the L synapses have zero strength. Output spikes that are generated at this instant will cause a number of synapses in the N population to change strength: the symmetric STDP rule (grey curve) specifies strengthening of synapses that become active nearby in time and therefore close to the wavefront, indicated by the upward pointing orange arrows in Fig. 6.1B. The STDP rule also specifies weakening of synapses that become active further away in time and therefore further away from the wavefront, indicated by the downward pointing orange arrows. Thus, at this instant, increases in synaptic strength in the N population align with synapses that are already strong, and decreases in synaptic strength align with synapses that are already weak. In the L population, the wavefront is some distance, equal to the wave speed multiplied by the population delay, behind the wavefront in the N population. Synapses at and around the wavefront in the L population will also be strengthened (upward pointing light blue arrows) because of the coincident activity in the output, even though the L wavefront doesn't actually drive the output during this snapshot. Similarly, the current output activity causes synapses in the L population that are distant from the wavefront to weaken (downward pointing light blue arrows). Thus, output spikes that are driven by activity in the N population causes synapses to strengthen at the left, and to weaken at the right, of peaks in the L connectivity

pattern. This will cause the L connectivity pattern to shift towards the left.

A snapshot of the wavefront at a later time is drawn in Fig. 6.1C. Here, the wavefront is further to the right, such that it is activity in the L population, and not the N population, that drives spiking in the output. At this instant, output spikes cause synapses to strengthen at the right, and to weaken at the left, of peaks in the N connectivity pattern (orange arrows), because the N wavefront is further ahead of the L wavefront. This will cause the N connectivity pattern to shift towards the right. In the L population, synaptic increases align with synapses that are already strong, and synaptic decreases align with weak synapses (light blue arrows).

In summary, wavefront activity in a given input population drives activity in the output neuron, the consequence of which is to induce a spatial shift in the connectivity pattern of the other population. As this process unfolds while more wavefronts traverse the inputs, the two connectivity patterns shift in opposite directions. In Fig. 6.1D, we examine a snapshot of the model much later on, when the connectivity patterns have moved apart. Now, when the wavefront in the N population lies at a peak in the N connectivity pattern, the wavefront in the L population lies at a peak in the L connectivity pattern. As such, the spatial profile of synaptic strengthening in the both the N and L populations overlaps the region of strong synapses in the respective N and L connectivity patterns (upward orange and blue arrows, respectively). Likewise, synaptic weakening in both populations occurs where the synapses are already weak (downward orange and blue arrows). As a result, the wavefront in the N population no longer induces a spatial shift in the L connectivity pattern. Likewise, the wavefront in the L population no longer induces a spatial shift in the N connectivity pattern.

So long as the population delay, t_d , STDP rule and wave speed do not change,




Figure 6.1: **Model framework and dynamics.** **A.** The model used in this chapter builds upon the one-dimensional model used in previous chapters by comprising multiple input populations, each with their own set of excitatory synapses that converge on a single output neuron. Wavefronts in each population are staggered in time by a fixed delay and travel from left to right. When recruited by a wave, each input neuron generates a burst of spikes from a Bernoulli process. **B–D.** Schematic for the development of connectivity patterns in a realisation of the model with two input populations. Each panel consists of a plot of the current connectivity patterns (dark coloured square waves), the wavefront location (dark coloured circle) in each input population (light coloured circles), the STDP rule (grey curve) and the synaptic changes specified by the STDP rule (coloured arrows) that result from activity in the input populations and a common output neuron. **B.** When the inter-population delay, t_d , is zero, the two connectivity patterns develop the same periodic connectivity pattern. If t_d is then set to a non-zero value, output activity that is driven by the wavefront in the N population induces synaptic changes around the wavefront in the L population, which lags behind to the left. This causes peaks in the L connectivity pattern to move leftwards. **C.** By the same reasoning as in **B**, output activity that is driven by the wavefront in the L population induces synaptic changes around the wavefront in the N population, which is in front to the right. This causes peaks in the N connectivity pattern to move rightwards. **D.** After many waves have passed, the N and L connectivity patterns have moved in opposite directions. Here, when the wavefront in the N population lies at a peak in the N connectivity pattern, the wavefront in the L population also lies at a peak in the L connectivity pattern. As such, synaptic changes due to activity in the output neuron only act to further strengthen strong synapses and to weaken weak synapses. Consequently, the connectivity patterns no longer shift in opposite directions, but have reached a stable configuration.

this configuration of spatially shifted connectivity patterns will not change with following wavefronts. Moreover, this configuration requires that the two connectivity patterns be shifted relative to each other by exactly the same distance that separates the wavefront in each input population, vt_d , which is the distance travelled by a wave with wave speed v during the time t_d . Such a process has been described previously in modelling studies of DS simple cell development (Wimbauer et al., 1997a; Blais et al., 2000), in which plasticity was governed by a rate based Hebbian rule. In the following section, we show that this is indeed the case for any STDP rule by deriving a general description for the synaptic strength dynamics during which wavefronts traverse two, delayed input populations.

Deriving synaptic strength dynamics for two input populations with a delay

In order to derive an expression for the synaptic strength dynamics, we pursue the same approach taken to derive Eq. 3.4 in Chapter 3, modified to incorporate both input populations and their respective synaptic strengths. Thus, we begin with a general expression for the dynamics of synaptic strengths connecting input population j (either N or L) to a single output neuron:

$$\partial_T w_j(x) = \eta \int_{-\infty}^{\infty} d\Delta t K(\Delta t) C_j(x, \Delta t). \quad (6.12)$$

As in Chapter 3, $K(\Delta t)$ is the STDP rule as a function of the input-output spike latency, Δt , and $C_j(x, \Delta t) = \int_{-\infty}^{\infty} dt S_j(t + \Delta t) S_{\text{out}}(t)$ is the time-averaged cross-correlation between activity in the input population j and the output neuron over the duration of a single wave, T . In order to compute $C_j(x, \Delta t)$, we require a description for these activity patterns. For simplicity, we omit excitatory postsynaptic potentials (EPSPs) and input burst waveforms for the time being. These enter the derivation in a simple way, as shown in Chapter 3, and do not substantially affect the main results of this chapter. The firing rate in the N and L input populations, $S_N(t)$ and

$S_L(t)$, can be expressed as two pulses travelling from left to right, with the L pulse lagging the N pulse:

$$S_N(x, t) = \delta(x - vt) \quad (6.13)$$

$$S_L(x, t) = \delta(x - v(t - t_d)). \quad (6.14)$$

The output firing rate, which is the linear sum of activation from both N and L populations, weighted by the synaptic strengths, $w_N(x)$ and $w_L(x)$, becomes

$$\begin{aligned} S_{\text{out}}(t) &= \int_{-\infty}^{\infty} dx w_N(x) \delta(x - vt) + w_L(x) \delta(x - v(t - t_d)) \\ &= w_N(vt) + w_L(vt - vt_d). \end{aligned} \quad (6.15)$$

Using Eqs. 6.13–6.15, the cross-correlation between each input population and the output neuron can be found thus:

$$C_N(x, \Delta t) = w_N(x - v\Delta t) + w_L(x - v\Delta t - vt_d) \quad (6.16)$$

$$C_L(x, \Delta t) = w_N(x - v\Delta t + vt_d) + w_L(x - v\Delta t), \quad (6.17)$$

and by substituting C_N and C_L into into Eq. 6.12, we obtain the following expressions for the synaptic strength dynamics:

$$\partial_T w_N(x) = \eta[K_v(x) * w_N(x) + K_v(x - vt_d) * w_L(x - vt_d)] \quad (6.18)$$

$$\partial_T w_L(x) = \eta[K_v(x + vt_d) * w_N(x + vt_d) + K_v(x) * w_L(x)], \quad (6.19)$$

where we have rescaled K such that $K_v(x) = v^{-1}K(x/v)$, as we had done in Chapter 3. Taking the Fourier transform of Eqs. 6.18 and 6.19, the dynamics can be written as a coupled partial differential equation in order to find solutions for w_N and w_L as $T \rightarrow \infty$. Thus,

$$\partial_T \begin{pmatrix} \tilde{w}_N(k) \\ \tilde{w}_L(k) \end{pmatrix} = \eta \tilde{K}_v(k) \begin{pmatrix} 1 & e^{-ivt_d k} \\ e^{ivt_d k} & 1 \end{pmatrix} \begin{pmatrix} \tilde{w}_N(k) \\ \tilde{w}_L(k) \end{pmatrix}, \quad (6.20)$$

where ‘ $\sim$ ’ denotes the Fourier transformed object and k is the spatial frequency. When the additional dynamics of the input firing rate, $\alpha(t)$, and EPSP, $\epsilon(t)$, are taken into account, they act equally on both populations by filtering $\partial_T w_j(x)$, in the same way as for the single population case described in Chapter 3. We can follow the same steps as in Chapter 3 and amalgamate the spatial projections of $\alpha(t)$ and $\epsilon(t)$, respectively $\alpha(x/v)$ and $\epsilon(x/v)$, by convolving them with $K_v(x)$. Recall from Chapter 3 that $\alpha(x/v)$ enters the dynamics twice, once as a convolution and once as a cross-correlation. Thus, we replace $K_v(x)$ in Eq. 6.20 with $\kappa(x) = K_v(x) * \alpha(-x/v) * \alpha(x/v) * \epsilon(x/v)$. The two eigenvalues of $\eta \tilde{\kappa}(k) \begin{pmatrix} 1 & e^{-ivt_d k} \\ e^{ivt_d k} & 1 \end{pmatrix}$ are therefore $\sigma_1 = 0$ and $\sigma_2 = 2\eta \tilde{\kappa}$, the latter taking the same form as the eigenvalue for the single population case derived in Chapter 3. In order to find time dependent solutions for w_N and w_L , we impose an initial condition of random synaptic strengths in both the N and L synapse populations such that $\tilde{w}_N|_{T=0} = \tilde{w}_L|_{T=0} = \tilde{w}_0$, where the expected form of $\tilde{w}_0$ is flat. As such,

$$\tilde{w}_N = \frac{\tilde{w}_0}{2} \left[1 - e^{-ivt_d k} + (1 + e^{-ivt_d k}) e^{2\eta \tilde{\kappa} T} \right] \quad (6.21)$$

$$\tilde{w}_L = \frac{\tilde{w}_0}{2} \left[1 - e^{ivt_d k} + (1 + e^{ivt_d k}) e^{2\eta \tilde{\kappa} T} \right]. \quad (6.22)$$

For $T \rightarrow \infty$, Eq.s 6.21 and 6.22 can be approximated by

$$\tilde{w}_N \simeq \frac{\tilde{w}_0}{2} (1 + e^{-ivt_d k}) e^{2\eta \tilde{\kappa} T} \quad (6.23)$$

$$\tilde{w}_L \simeq \frac{\tilde{w}_0}{2} (1 + e^{ivt_d k}) e^{2\eta \tilde{\kappa} T}. \quad (6.24)$$

For the STDP rules used here, $\tilde{\kappa}(k)$ has a positive valued maximum at a spatial frequency, k^* . Thus, as with the case for a single input population, the eigenmode corresponding to k^* grows exponentially faster than all other modes and dominates the connectivity patterns for both N and L populations. In other words, both w_N and w_L become periodic in space with the same spatial frequency, k^* . However, we

also find that $\tilde{w}_L$ is a phase shifted copy of $\tilde{w}_N$:

$$\tilde{w}_L = \tilde{w}_N e^{i v t_d k}. \quad (6.25)$$

This phase shift occurs for any STDP rule and, for the STDP rules considered here, under the influence of which a single spatial frequency becomes dominant as $T \rightarrow \infty$, the spatial phase shift can be expressed:

$$\Delta\phi = v t_d k^*. \quad (6.26)$$

The positive phase shift indicates that, relative to w_N , w_L has shifted in a direction opposite to that of the travelling wave. Furthermore, the relative shift in space between $w_N|_{T \rightarrow \infty}$ and $w_L|_{T \rightarrow \infty}$ is $\Delta\phi/k^* = v t_d$, the distance travelled by the wave during the relative inter-population delay. It is important to note that $k^* \propto 1/v$, as $\kappa(x)$ is scaled in space by the wave speed. Thus, $\Delta\phi$ does not, in fact, vary with the wave speed. Rather, by writing $\kappa(x) = \kappa(vt)$, where $\kappa(vt)$ has a characteristic temporal frequency of $f^* = v k^*$, we have $\Delta\phi = t_d f^*$.

We may make an even more general statement here. In Fig. 3.6 in Chapter 3, we found that the pattern of wave initiation had little effect on k^* . Only when waves were generated with strictly regular intervals at a sufficiently high frequency, was k^* modified to higher frequencies. In other words, $\Delta\phi$ depends very little on the stimulus pattern, but depends heavily on the population delay and the temporal dependencies of $K(\Delta t)$, $\alpha(t)$ and $\epsilon(t)$. Thus, regardless of the speed with which visual features move within the visual field, the phase offset can be built into the visual circuitry with appropriately set time delays, STDP rules and temporal dynamics of individual neurons. This may be a useful property for the development of DS simple cells in V1.

A leading model for DS simple cell RFs sums input from two, non-DS Gabor RFs, one of which is delayed and has a spatial phase offset with respect to the other. The application of our results to this particular model of direction selectivity therefore relies on the assumption that the characteristic spatial frequency of Gabor RFs is also established by the periodic patterning of feedforward circuits as outlined in this thesis. We begin to address the applicability of these results to the development of DS simple cells later in this chapter. First, we seek to verify our analytical results using a simulated network of spiking neurons in the following section.

Development of connectivity patterns from two delayed input populations in a simulation of spiking neurons

In order to verify the analytical results above in a more biologically realistic model, we extended our simulation of travelling waves in a one-dimensional (1D) network of 500 spiking neurons, used in previous chapters, to model the developing connectivity between a common output neuron and two input populations, a non-lagged population, N, and a lagged population, L. In each simulation, 2000 waves in total drove plasticity in the synaptic strengths, whereby waves traversed the L population with a temporal delay of t_d seconds relative to waves in the N population, as depicted in Fig. 6.1A. In both N and L populations, wavefronts travelled with a speed of $50^\circ/\text{s}$ in the rightward direction only, as bidirectional waves would have resulted in a cancelling of spatial offsets in the development of the connectivity patterns. Full details of the simulation are provided in the Methods.

Delays in thalamocortical circuits that have been proposed to play a role in direction selectivity range from tens of milliseconds (Baker & Bair, 2012) to hundreds of milliseconds (Saul & Humphrey, 1990). For now, we avoid constraining our analytical results to any particular model of delay times, and examine their robustness over a wide range. To do this, we conducted a set of simulations, between which t_d

was varied in the range 0–475 ms. Simulations were conducted for both asymmetric and symmetric STDP rules (see General Methods), using a fixed decay time constant of $\tau_+ = 20$ ms. For each t_d , we repeated 16 trials of the simulation using different seeds so as to generate unique spike sequences. In Fig. 6.2A (using the asymmetric STDP rule) and B (using the symmetric STDP rule), we provide examples of N (red) and L (blue) connectivity patterns that developed in these simulations with different values of t_d (Fig. 6.2A, top to bottom: $t_d = 0, 75, 150, 225, 300$ and 375 ms; Fig. 6.2B, top to bottom: $t_d = 0, 50, 100, 150, 200$ and 250 ms). For both STDP rules, the N and L populations established nearly identical periodic connectivity patterns when $t_d = 0$ ms. As t_d was increased, the phase of the L connectivity pattern increased with respect to that of the N connectivity pattern.

We first ensured that the dominant spatial frequency of the connectivity patterns, k_{sim}^* , did not vary with the population delay (t -test for probability that $g = 0$ where $k_{\text{sim}}^* = gt_d$: $p > 0.3$ with the asymmetric STDP rule and $p > 0.1$ with the symmetric STDP rule), where the subscript ‘sim’ distinguishes values measured from each simulation from corresponding values that were predicted numerically. However, there was a small but significant difference between the measured, k_{sim}^* , and predicted, k^* , dominant spatial frequencies (asymmetric STDP: $k_{\text{sim}}^* = 0.0640 \pm 0.0002$ cycles/ $^\circ$, $k^* = 0.076$ cycles/ $^\circ$, t -test with $p < 0.001$; symmetric STDP: $k_{\text{sim}}^* = 0.0938 \pm 0.0003$ cycles/ $^\circ$, $k^* = 0.1000$ cycles/ $^\circ$, t -test with $p < 0.001$; mean $\pm$ SEM). Possible reasons for this is the non-linear hard bounds applied to the synaptic strengths to keep them in the range 0–1, boundary effects due to the finite size of the input layer, or the finite resolution with which neurons sample the input space. Further tests are needed to distinguish these possible causes. Nevertheless, as we are about to show, this did not affect the dependence of the phase shift on t_d .

We quantified the observed phase shift by computing the cross-correlogram (CCG,




Figure 6.2: **Development of connectivity patterns in simulations with two input populations.** **A & B.** Examples of the connectivity patterns that developed in simulations with different population delays. Red and blue patterns correspond to w_N and w_L , respectively. We plot w_L with a vertical offset for clarity. In the simulation, w_L has the same hard bounds as does w_N . Only the central 180° of the connectivity patterns are shown. Dark shaded lines provide a visual guide between peaks in the connectivity pattern that are separated by a distance vt_d . **A.** Simulations with an asymmetric STDP rule. Top to bottom: $t_d = 0, 75, 150, 225, 300$ and 375 ms. **B.** Simulations with a symmetric STDP rule. Top to bottom: $t_d = 0, 50, 100, 150, 200$ and 250 ms. **C.** CCGs, averaged over all 16 trials, between the two connectivity patterns in simulations with t_d values of 0, 200 and 425 ms and with an asymmetric STDP rule. Spatial shifts between the connectivity patterns, $D_{\text{sim}}(w_N, w_L)$, were extracted from the location of the dominant peak in the CCG. **D** and **E.** Plots of the measured phase shifts, $\Delta\phi_{\text{sim}} = D_{\text{sim}}(w_N, w_L) \bar{k}_{\text{sim}}^*$, as a function of t_d (black markers, mean $\pm$ SEM), for simulations with asymmetric (**D**) and symmetric (**E**) STDP rules. In blue are the predicted phase shifts, $\Delta\phi = vt_d \bar{k}_{\text{sim}}^*$.

see Methods) between the N and L connectivity patterns from each simulation. Three such CCGs, averaged over all trials, are shown in Fig. 6.2C for $t_d = 0, 200$ and 425 ms (black, blue and green curves, respectively). CCGs exhibited a single, dominant peak at the location, $D_{\text{sim}}(w_N, w_L)$, which reflects the inter-dependence of $w_N(x)$ and $w_L(x - vt_d)$ during development. We converted the spatial offset of the CCG peak to a phase shift by multiplying $D_{\text{sim}}(w_N, w_L)$ with the dominant spatial frequency, k_{sim}^* . In Fig. 6.2D and E, we plot $\Delta\phi_{\text{sim}} = D_{\text{sim}}(w_N, w_L)k_{\text{sim}}^*$ as a function of t_d . Against these measurements, we plot the predicted phase difference using $\Delta\phi = vt_d\bar{k}_{\text{sim}}^*$, where $v = 50^\circ/\text{s}$ and we have used the dominant spatial frequency averaged over all trials and all t_d values, $\bar{k}_{\text{sim}}^*$, instead of the theoretical dominant spatial frequency, k^* , which differed slightly. Our model provides an accurate description of the relationship between t_d and the phase offset between the two connectivity patterns, with coefficients of determination between the measured and the predicted phase shifts of $R^2 > 0.98$ for both the asymmetric and symmetric STDP rules.

Synaptic strength dynamics for multiple delayed input populations

A common feature of response properties in the visual system is a distribution of response latencies to visual stimuli, which can range from tens to hundreds of milliseconds at all stages along the visual pathway (Daniels et al., 1978; Blakemore & Vital-Durand, 1986a; DeAngelis et al., 1993a; Tavazoie & Reid, 2000). In order to extend our analytical results to an arbitrary number of input populations with different response latencies, we generalise our analytical derivation for two input populations by describing wave activity in population m with a response latency of t_m :

$$S_m(x, t) = \delta(x - v(t - t_m)). \quad (6.27)$$

The firing rate of the output neuron is proportional to the linear sum of synaptic inputs from all populations, weighted by their respective synaptic strengths:

$$\begin{aligned} S_{\text{out}}(t) &= \sum_m \int_{-\infty}^{\infty} dx S_m(x, t) w_m(x) \\ &= \sum_m w_m(v(t - t_m)). \end{aligned} \quad (6.28)$$

Thus, the cross-correlation between input population m and the output neuron becomes

$$\begin{aligned} C_m(x, \Delta t) &= \int_{-\infty}^{\infty} dt S_m(x, t) S_{\text{out}}(t) \\ &= \int_{-\infty}^{\infty} dt \delta(x - v(t + \Delta t - t_m)) \sum_n w_n(v(t - t_n)) \\ &= \sum_n w_n(x - v(\Delta t + t_{mn})), \end{aligned} \quad (6.29)$$

where $t_{mn} = t_m - t_n$ is the difference in response latencies between input populations m and n . The synaptic strength dynamics for population m can be found by substituting Eq. 6.29 into Eq. 6.12. Incorporating the input firing rate waveform and the EPSP, and taking the Fourier transform, we have

$$\partial_T \tilde{w}_m(k) = \eta \tilde{\kappa}(k) \sum_n e^{i v t_{mn} k} \tilde{w}_n(k). \quad (6.30)$$

Note that, for two input populations, Eq. 6.30 and Eq. 6.20 are equivalent. Indeed, as $T \rightarrow \infty$, the phase difference between any two connectivity patterns, belonging to populations m and n , must be proportional to the relative delay between those populations: $\Delta\phi(w_m, w_n) = v t_{mn} k^*$. We build upon this result to address a potential role for different response latencies in the LGN in the development of direction selective geniculocortical circuits.

Development of connectivity patterns from four delayed input populations in a simulation of spiking neurons

In this section, we examine whether the delay between two input populations determines the phase shift between their respective connectivity patterns in simulations of spiking neurons with more than two input populations. We do so in the context of a developing geniculocortical circuit, which is a leading model for the development of direction selectivity in V1. This model proposes that, in addition to spatially offset non-lagged and lagged inputs, simple cells receive spatially offset inputs from two other cell classes in the LGN: ON cells, which respond to local increases in luminance, and OFF cells, which respond to local decreases. ON and OFF cells respond well to changes in luminance, which are most often caused when objects move into and out of their RFs. Wavefronts of spiking activity reliably track moving edge stimuli across populations of ON and OFF cells in the primate retina (Frechette et al., 2005), and it is plausible that such activity patterns are relayed in the LGN. Thus, an ON and an OFF cell with RFs that overlap in space will respond at different times as the object moves into the RF or as it exits the RF, which corresponds to changes in luminance with opposite signs.

We therefore simulated the development of connectivity patterns under the influence of travelling waves in a circuit with four input populations. To each population, we assigned a response characteristic from the LGN: ON and non-lagged (with synaptic strengths w_{ON}^{N}), ON and lagged (w_{ON}^{L}), OFF and non-lagged ($w_{\text{OFF}}^{\text{N}}$), and OFF and lagged ($w_{\text{OFF}}^{\text{L}}$). We imagined travelling waves in the input layer being driven by a bright spot stimulus moving left to right on a dark background, which we refer to as the training stimulus. Thus, at a given location on the input layer, ON cells responded prior to OFF cells within the same latency group (N or L). We ignore activity due to the opposite polarity RF surround, which is weaker in young animals

than in adults (Cai et al., 1997). For simplicity, we predetermine the order in which each population responds: ON non-lagged (ON-N) followed by ON lagged (ON-L), OFF non-lagged (OFF-N) and finally OFF lagged (OFF-L). We also fix the relative time delay, t_d , between successive populations, such that $t_{mn} = (m - n)t_d$ in Eq. 6.30, where the integers m and n indicate the two populations by ranking their response latencies, with $m = 1$ denoting ON-N and so on. We varied t_d between simulations in the range 0–150 ms, which meant that OFF-L inputs, which had the longest response latency, were delayed by a maximum of 450 ms relative to ON-N inputs.

In Fig. 6.3A and B (asymmetric and symmetric rules, respectively), we provide examples of connectivity patterns that developed in the simulations with different values of t_d (Fig. 6.3A: $t_d = 0, 75$ and 150 ms, top to bottom; Fig. 6.3B: $t_d = 0, 50$ and 100 ms, top to bottom). We quantified the phase shift between the ON-N connectivity pattern and the three remaining connectivity patterns using the same method as for the case with two populations above. In Figs. 6.3C and D, we plot the measured phase differences, $\Delta\phi_{\text{sim}}(w_{\text{ON}}^{\text{N}}, w_{\text{ON}}^{\text{L}})$ (orange markers), $\Delta\phi_{\text{sim}}(w_{\text{ON}}^{\text{N}}, w_{\text{OFF}}^{\text{N}})$ (dark blue markers) and $\Delta\phi_{\text{sim}}(w_{\text{ON}}^{\text{N}}, w_{\text{OFF}}^{\text{L}})$ (light blue markers), as a function of the relative population delay, t_{mn} , where $\Delta\phi_{\text{sim}}(w_{\text{ON}}^{\text{N}}, w_{\text{A}}^{\text{B}}) = D_{\text{sim}}(w_{\text{ON}}^{\text{N}}, w_{\text{A}}^{\text{B}})k_{\text{sim}}^*$. Against these, we plot the predicted phase differences (grey lines), which match the measured values with coefficients of determination of $R^2 > 0.99$. As in the case with two input populations, we found that the spatial frequencies of the connectivity patterns differed by a small but significant amount from the predicted spatial frequencies ($\sim 3\text{--}5\%$ difference between measured and predicted spatial frequencies, $p < 0.001$ for all population types and STDP rules, t -test), but did not vary significantly with t_{mn} (t -test probability that $g = 0$ where $k_{\text{sim}}^* = gt_{mn}$: $p > 0.21$ for all populations and STDP rules). In conclusion, our simulation results support the theoretical result reached in Eq. 6.30, that even with multiple input populations, the phase shift




Figure 6.3: **Development of connectivity patterns in simulations with four input populations.** **A & B.** Examples of the connectivity patterns that developed in simulations with different population delays. Dark red: w_{ON}^N ; orange: w_{ON}^L ; dark blue: w_{OFF}^N ; light blue: w_{OFF}^L . Populations with longer delays, $t_{mn} = (m - n)t_d$, are drawn with larger vertical offsets for clarity. Only the central 180° of the connectivity patterns are shown. Dark shaded lines provide a visual guide between peaks in the connectivity pattern that are separated by a distance vt_{mn} . **A.** Simulations with an asymmetric STDP rule. Top to bottom: $t_d = 0, 75$ and 150 ms. **B.** Simulations with a symmetric STDP rule. Top to bottom: $t_d = 0, 50$ and 100 ms. **C and D.** Plots of the measured phase shifts, $\Delta\phi_{sim} = D_{sim}(w_{ON}^N, w_A^B)k_{sim}^*$, as a function of t_d , where A indicates ON or OFF and B indicates N or L. Markers provide mean $\pm$ SEM. Grey lines are the predicted phase shifts, $\Delta\phi = vt_{mn}k_{sim}^*$. **C.** Simulations use an asymmetric STDP rule. **E.** Simulations use a symmetric STDP rule.

between any two populations is determined by the relative delay in the responses of those two populations.

In summary of the results thus far, we have demonstrated that travelling waves in multiple, delayed input populations can act through STDP to build phase shifts between their respective connectivity patterns. In particular, phase shifts are independent of the speed of the wave, which, in many cases, implies that offsets can be determined by properties intrinsic to the neural circuit. This may prove particularly useful during the development of DS simple cells, as relative temporal delays due to intrinsic neuron or network properties change by only a small amount during the initial seeding of direction selectivity (Blakemore & Vital-Durand, 1986a; Cai et al., 1997; Tavazoie & Reid, 2000), thus predetermining spatial offsets between inputs with a given delay. However, the relative timing of ON and OFF activity depends largely on the spatiotemporal structure of natural scenes, which can be considerably varied. We address this issue in the Discussion. In the following section, we explore how the development of spatially and temporally offset connectivity patterns, by the mechanism shown above, may be applied to the development of direction selective properties in simple cells.

Dependence of direction selectivity on the relative time delays between inputs

In this section, our goal is to lay groundwork for understanding how the mechanism examined above relates to the development of direction selectivity of simple cells. We examine how the different arrangements of ON-N, ON-L, OFF-N and OFF-L inputs, which developed in the above simulations, influence the direction selectivity of a receiving output neuron, and whether other configurations that are not permitted by this developmental scheme might yield greater direction selectivity. An examination of the development of more realistic simple cell RFs, which are well described by two dimensional Gabor filters (Jones & Palmer, 1987), is left for future

work.

To assess the response properties of the output neuron that developed in the above simulations, we constructed linear spatiotemporal RFs (STRFs) by multiplying the connectivity patterns, $w_A^B(x)$, with a stimulus response function of synaptic input to the output neuron, $h_A^B(t)$ (Fig. 6.4A), where A denotes ON or OFF and B denotes N or L, and summing these spatiotemporal contributions across the four populations (see Methods). In this case, the stimulus response is proportional to the convolution in time of the input firing rate, $\alpha(t)$, with the EPSP, $\epsilon(t)$. In Fig. 6.4B, we show examples of STRFs that are taken from simulations in which t_d , the time lag between successively delayed populations, took values of 25–100 ms. Note that t_d determined the delay between ON and OFF population types as well as between N and L population types. Thus, as t_d was increased, the ON- and OFF-type inputs as well as the N- and L-type inputs would develop greater shifts in space, while the L-type inputs were also delayed by longer times with respect to the N-type inputs. When $t_d = 50$ or 75 ms (Fig. 6.4B, central panels), the distinct tilt in the STRFs that characterises DS simple cells could be seen.

To quantify the direction selectivity of STRFs that developed in these simulations, we examined the degree to which the firing rate of the output neuron was modulated by a test stimulus, which comprised a drifting sinusoidal luminance grating with a range of possible spatial and temporal frequencies. The firing rate of a linear output neuron can be computed by convolving the STRF with the stimulus, which, in frequency space, is equivalent to computing the Fourier transform of the STRF and multiplying it with a delta function at the spatial and temporal frequency of the stimulus. The modulation magnitude, for all spatial and temporal frequencies, is therefore equal to the amplitude spectrum of the STRF, and reports the sensitivity of the STRF to a sinusoidal grating moving in a particular direction as a function

of the spatial and temporal frequency of the grating (Adelson & Bergen, 1985). An example of an amplitude spectrum is given in Fig. 6.5A for a simulation in which an asymmetric STDP rule was used and $t_d = 100$ ms. STRFs that are direction selective yield power spectra that are asymmetric about the origin, with a dominant peak in two quadrants, and a subdominant peak in the remaining two quadrants. In this case, the dominant peak is in the upper right quadrant (hereafter referred to as quadrant 1) and lower left quadrant, indicating a preference for motion in the rightwards direction. The peak in the upper left quadrant (quadrant 2) reveals that the STRF is also sensitive to motion in the leftwards direction, albeit to a lesser extent. A direction selectivity index (DSI) of the STRF can be defined as the difference in magnitude between these two peaks ($r_1 - r_2$ in Fig. 6.5A), normalised by the sum of their magnitudes (Eq. 6.10 in Methods). The DSI for STRFs that developed in our simulations is plotted as a function of t_d in Fig. 6.5B, for simulations using both asymmetric and symmetric STDP rules. Naturally, direction selectivity was absent when $t_d = 0$ ms, due to the absence of any spatial or temporal offsets between the four input populations. The DSI peaked when t_d was between 75–100 ms, and dropped to negative values as t_d increased further, demonstrating a shift in direction selectivity from motion in the same direction as the training stimulus to motion in the opposite direction for $t_d \gtrsim 125$ ms. The DSIs computed for these STRFs are relatively small compared to those observed in V1 of the ferret, cat and monkey (Saul & Feidler, 2002; Gur et al., 2005; Moore et al., 2005). There are several likely reasons for this, which we review in the Discussion.

Given the STDP rules and stimulus response characteristics used in our simulations, do temporal delays of 75–100 ms yield STRFs that have the highest possible DSI, or do other STRF configurations provide greater DSIs? It is difficult to interpret from these simulations how t_d influences the direction selectivity of STRFs in the way shown in Fig. 6.5B, as the time delay between the four input populations

affects three STRF features: the spatial offset between ON and OFF populations, the spatial offset between N and L populations, and the temporal offset between N and L populations. In order to untangle the effect these features have on direction selectivity, we constructed a new STRF, which was modelled on those that developed in the simulations, and manually varied its spatial and temporal structure in these three, independent ways. We refer to this new STRF, and its spatial and temporal permutations, as the artificial STRF, to distinguish it from STRFs obtained from the connectivity patterns that developed in the simulations. Fig. 6.5C shows examples of different temporal offsets between N and L components (left column), spatial offsets between N and L components (middle column), and spatial offsets between ON and OFF components (right column) of the artificial STRF. For each configuration of the artificial STRF, we measured the DSI from its amplitude spectrum. First of all, we point out that spatial offsets between ON and OFF components modulate the response amplitude of the model output neuron for test stimuli moving in both directions, but do not affect direction selectivity, as they have no effect on the tilt of the STRF in space and time. We therefore concentrate on the spatial and temporal offsets between the N and L components and their effect on the DSI. We computed the DSI for each permutation of the artificial STRF, whereby the N-L spatial shift was moved through one complete cycle of the periodic pattern and the N-L delay was increased through 0–200 ms. Fig. 6.5D plots the DSI as a function of these two variables. In green and yellow are the predicted phase shifts, $\Delta\phi = vt_d\bar{k}_{\text{sim}}^*$, as a function of t_d , for the asymmetric and symmetric STDP rules, respectively. The DSI along these lines approximately corresponds to the DSI measured from the simulation STRFs in Fig. 6.5B.

The blue boundary in Fig. 6.5D encircles DSIs that exceed 90% of the maximum DSI observed over all configurations of the artificial STRF. It is apparent that the mechanism by which STDP instructs the development of direction selectivity does

not produce STRFs that are maximally direction selective in our simulations. What changes to our simulation could yield STRFs that are more direction selective, i.e. what variables could increase the slope of $\Delta\phi$ (green and yellow lines) in Fig. 6.5D? Recall from above that the relationship between $\Delta\phi$ and t_d does not depend on the speed of the training stimulus, as faster wave speeds give rise to connectivity patterns with lower spatial frequencies. Alternatively, the stimulus response functions could be made faster. However, this would also have the effect of compressing the map of DSI values along the N-L delay axis, moving the blue region to even smaller N-L delays. Thus, in order for maximally direction selective STRFs to develop, the spatial frequency of the connectivity pattern must be increased by decreasing the STDP time constant. The STDP time scales needed to move the curve of $\Delta\phi$ into the blue region would be $\sim 1\text{--}3\text{ ms}$ for the asymmetric rule and $\sim 3\text{--}10\text{ ms}$ for the symmetric rule, considerably shorter than STDP time scales typically observed experimentally (Caporale & Dan, 2008).

Direction selectivity alone may not be an appropriate feature to maximise, as it is feasible for STRFs to yield large DSIs but have relatively small response amplitudes. Signals from a neuron with a weak response amplitude might be swamped in a population of neurons with stronger, but less selective responses. In Fig. 6.5E, we characterise the maximum response amplitude (MR) of the artificial STRF as a function of the N-L shift and N-L delay, which shows that the largest MRs are obtained with STRF configurations that yield relatively low DSIs. Thus, a more useful property of a DS cell may be to jointly maximise its DSI and MR. In Fig. 6.5F, we plot the product of the DSI and MR for the artificial STRF as a function of the N-L shift and N-L delay. Outlined in blue are artificial STRF configurations that yield high DSIs alone, as in Fig. 6.5D, yet which have MRs that are only 40–70% of the largest observed MR. Outlined in purple are STRF configurations that yield joint high DSIs and large MRs. The minimum DSI for these is 78% of the highest




Figure 6.4: **Spatiotemporal receptive fields (STRFs) in simulations with four input populations.** **A.** Stimulus response functions are the convolution in time of the input firing rate response, $\alpha(t)$, with the EPSP, $\epsilon(t)$. They describe the synaptic activation of the output neuron from inputs in each of the four populations, under the assumption that travelling waves were induced by a training stimulus consisting of a moving bright bar against a dark background. ON input types have positive stimulus responses, whereas OFF input types have negative responses, so that decreases in luminance elicit excitatory activation of the output neuron. **B.** Examples of STRFs from simulations with different inter-population delays, t_d . Top row: simulations with an asymmetric STDP rule; bottom row: simulations with a symmetric STDP rule. STRFs were constructed by summing the products of each connectivity pattern with its corresponding stimulus response function. Only the central 60° of the STRFs are shown.

observed DSI. However, their MRs are between 70–90% of the largest observed MR. Coinciding with the artificial STRF configurations that are joint maximal are those STRFs that yielded the greatest DSI scores in our spiking neuron simulations, for which $t_d = 75\text{--}100\text{ ms}$ (intersection of the green and yellow lines with the purple region in Fig. 6.5F). These results demonstrate that STDP rules with time scales of $\sim 20\text{ ms}$ may be appropriately tuned to build DS cells with relatively strong direction selectivity and response amplitudes.

Discussion

We have outlined a new approach to investigate a leading model for the development of orientation and direction selectivity in V1 simple cells, in which the spatial and temporal organisation of simple cell STRFs are seeded by the development of connectivity patterns with non-lagged and lagged ON and OFF cell types in the LGN. By considering the latencies of these cell types in response to a moving luminance contour, we have derived a general description for the influence of STDP on the development of multiple connectivity patterns between delayed input populations and a common output. In particular, we found that, regardless of the number of input populations, these connectivity patterns would develop the same periodic structure as in the case with a single input population, and that two connectivity patterns would develop a phase shift with respect to each other, by an amount proportional to the temporal delay between their respective inputs. Remarkably, we found that the phase shift depends heavily on properties that are intrinsic to the network, and is largely independent of the input activity pattern, beyond the requirement that it comprise travelling waves. In the following, we continue discussion of these general findings in the context of developing simple cell STRFs.




Figure 6.5: **Direction selectivity of STRFs from simulations with four input populations.** **A.** Amplitude spectrum of a STRF constructed from connectivity patterns that developed in a simulation with $t_d = 100$ ms and an asymmetric STDP rule. Amplitudes are normalised to the maximum amplitude, r_1 , which corresponds to the largest response amplitude of the output neuron to a luminance-modulated sinusoidal grating drifting in the rightwards direction. The largest response amplitude for motion in the leftwards direction corresponds to r_2 . The location of the peak amplitudes, r_1 and r_2 , are marked by the purple circles. **B.** Direction selectivity indices (DSIs) of STRFs as a function of t_d (mean $\pm$ SEM). Black markers: for simulations with an asymmetric STDP rule; grey markers: for simulations with a symmetric STDP rule. **C.** Artificial STRFs that have been modelled on those obtained from the simulations. Each column demonstrates two instances of the artificial STRF, either with different N-L temporal delays (left), N-L spatial shifts (middle) or ON-OFF spatial shifts (right). **D** DSI of artificial STRFs as a function of the N-L delay and N-L shift. Positive DSIs correspond to a preference for motion in the same direction as the training stimulus in the simulations. DSIs greater than 90% of the maximum DSI obtained are encircled in blue. Green and yellow lines plot the predicted phase shift between connectivity patterns in the simulation as a function of the N-L delay, t_d , for simulations with an asymmetric and symmetric STDP rule, respectively. **E.** Maximum response amplitude (r_1) in artificial STRFs as a function of the N-L delay and N-L shift, normalised to the maximum response amplitude obtained. **F.** Product of the DSI and maximum response amplitudes for model STRFs as a function of the N-L delay and N-L shift. Joint scores greater than 90% of the maximum obtained are encircled in purple. Blue, yellow and green curves are as in **D**.

Model limitations

We have chosen to model a linear circuit, with the exception of the bounded synaptic strengths, comprising only a single output neuron, so as to obtain tractable expressions for the synaptic strength dynamics (Eqs. 6.20 and 6.30) and derive analytic solutions (Eqs. 6.25–6.26). This is a considerable simplification of the simple cell circuitry. Simple cells receive a large fraction of cortical as well as thalamic input (Angelucci et al., 2002), and it is often the case that a linear STRF falls short of an accurate description of simple cell responses to visual stimuli, which can exhibit supralinear responses to motion in the preferred direction (Albrecht & Geisler, 1991; DeAngelis et al., 1993b), sublinear responses to motion in the anti-preferred direction (Reid et al., 1987; Reid et al., 1991; Rust et al., 2005), modulation from the surrounding extra classical RF (Blakemore & Tobin, 1972; DeAngelis et al., 1994), and non-linearities in spike generation (Priebe & Ferster, 2005). Whether, and by how much, network effects and non-linearities will influence the development of feedforward connectivity patterns by travelling wave activity and STDP is not clear. The robustness of our results to several sources of noise, described in Chapter 3, provides some confidence that the qualitative results in this chapter, regarding the development of periodic connectivity patterns with phase shifts between delayed input populations, may continue to hold true. Nevertheless, our results provide a well controlled foundation with which to refer when considering models of greater complexity.

Low direction selectivity indices of model spatiotemporal receptive fields

In order to examine direction selectivity, we constructed STRFs from the four connectivity patterns that developed in simulations. Unlike DSIs reported for DS simple cells in experimental studies, which often take values between 0.5–1, the DSIs associated with our model STRFs were no larger than 0.33. We are hesitant to infer from

this that a model of direction selectivity that is built on inputs from the LGN is a poor model, and suggest several reasons for this discrepancy. A significant fraction of simple cells are known to have linear STRFs that yield low DSIs, when compared to DSIs computed from the measured responses of the cell to motion stimuli (Reid et al., 1991; DeAngelis et al., 1993b). Nevertheless, many simple cells are found to have DSIs greater than 0.5, when determined by the linear STRF. Our method of calculating the DSI (Eq. 6.10) may also contribute to its low score in our models. The response amplitudes for preferred and anti-preferred motion are taken over all spatial and temporal frequencies from quadrants 1 and 2 in the amplitude spectrum of the STRF. To compute the DSI, we take the response amplitudes for motion in the preferred and anti-preferred directions to be the maximum amplitudes, over all spatial and temporal frequencies, in quadrants 1 and 2 of the STRF amplitude spectrum. However, DSIs are typically calculated at the spatial and temporal frequencies that provide the maximum response in the preferred direction only, which often do not produce the maximum response in the anti-preferred direction. In this case, the DSIs we report are necessarily lower than those reported elsewhere. A final reason for the low DSI scores here is that our model STRFs are built from inputs with only four distinct combinations of spatial and temporal offsets. However, cells in the LGN exhibit a continuous spectrum of RF locations and temporal delays, which may help yield a better defined space-time tilt in the STRFs of DS simple cells that receive their input.

Comparison with other models of simple cell development

A leading model for the development of orientation selective simple cells utilises a Hebbian plasticity rule to drive competition between negatively correlated ON and OFF inputs, which segregate in retinotopic space (Miller, 1994). However, the development of direction selectivity by the same plasticity rule relies on positive correlations between inputs that have different temporal responses to the same

stimulus, and thus do not compete. In this case, linear models of Hebbian plasticity fail to develop DS cells when the stimulus statistics contain no bias for motion in a particular direction (Wimbauer et al., 1997b; Wimbauer et al., 1997a; Blais et al., 2000). This is easily resolved by including a non-linearity in the dynamics of the model (Feidler et al., 1997; Blais et al., 2000). The linearity of the model we have presented here depends on the context of the input response properties. If we assign to each input an ON or OFF property, we must assume that their firing rates respond to changes in luminance with a non-linear threshold at zero. Although we have not explored it here, this non-linear response property should endow our model with the ability to break the symmetry of unbiased motion and develop selectivity for one particular direction.

One issue with models of the kind used in Wimbauer et al. (1997a) and Blais et al. (2000) is that the spatial frequency of STRFs develops independently of motion-related correlations in the input, as demonstrated in models that tackled orientation selectivity alone (Linsker, 1986b; Miller, 1994). Thus, the development of the spatiotemporal tilt in STRFs, due to correlations between differentially lagged inputs, is independent of the spatial structure. When the number of different temporal responses in the inputs is low, as in the Wimbauer et al. (1997a) and Blais et al. (2000) models, for which there are only non-lagged and lagged inputs, the separation between those developmental processes responsible for the spatial structure, and those responsible for the spatiotemporal tilt, imposes a restriction on the range of motion speeds that can yield strong direction selectivity (Blais et al., 2000). In particular, direction selectivity emerges only if the speed of the training stimuli, multiplied by the delay between the inputs, approximates a quarter cycle of the STRF's spatial profile. This may not in fact be a problem for geniculocortical circuits, due to the spectrum of response delays in the LGN (Daniels et al., 1978; Blakemore & Vital-Durand, 1986a; Tavazoie & Reid, 2000). Nevertheless, the model we have pre-

sented circumvents this issue, as both the spatial frequency and the spatiotemporal tilt of STRFs are determined by the motion of the training stimulus. As such, the phase shift between non-lagged and lagged inputs has very little dependence on the velocity of the training stimulus, regardless of the number of input neuron types. In fact, the results of this chapter and those in Chapter 3 suggest that the development of direction selectivity in our model has very little dependence at all on the pattern of the training stimulus, besides the requirement that it invoke travelling waves in the input layer.

Our model also contrasts substantially with previous models of STDP that address the development of direction selectivity. A widely explored feature of temporally asymmetric STDP rules is their ability to translate the spatial distribution of excitatory synaptic strengths in a direction opposite to the training stimulus (c.f. Fig. 3.6 in Chapter 3). Exploiting this to establish an imbalance between excitatory and inhibitory inputs can lead to direction selectivity in the postsynaptic cell (Rao & Sejnowski, 2003; Shon et al., 2004; Wensch et al., 2005; Honda et al., 2011). Experimental evidence for spatial shifts in neuronal RFs, after exposure to unidirectional stimuli, has been observed in the *Xenopus* tectum (Richards et al., 2010) and hippocampal place fields of rats (Mehta et al., 2000). However, the network structures that develop in these modelling studies should not give rise to the sharp spatial frequency tuning that is typical of simple cells. Buchs & Senn (2002) utilised the asymmetric STDP rule in a similar way, to shift the spatial distribution of depressing synapses, which also provide the necessary temporal offset (Chance et al., 1998), with respect to non-depressing synapses. Yet the success of their model in exhibiting direction selectivity hinges on hardwired constraints on the retinotopic location of depressing and non-depressing synapses, and depends also on a training stimulus that is periodic in time. In our own model, direction selectivity does not arise from this kind of spatial shift, which occurs in equal proportion for each

of the four connectivity patterns (all periodic connectivity patterns slowly migrate leftwards during training with rightward moving stimuli, as shown in Chapter 3). Rather, direction selectivity arises due the relative spatial shift between the connectivity patterns, which is a function of the difference in response latencies between the input populations. Moreover, it does not depend on the asymmetry of the STDP rule, but will work for any bandpass STDP rule. Thus, by the process described in this chapter, a DS cell might not require any hardwiring of the spatial organisation of its afferents, beyond the spatial restriction of its dendritic arbour, and the periodicity of the connectivity pattern provides the spatial frequency tuning that is characteristic of simple cells. However, a number of details first need to be worked out in our model, before we can make strong predictions about simple cell development. Several planned extensions to our model are described in the final section of this Discussion.

Organisation of ON and OFF inputs

In Miller (1994), Wimbauer et al. (1997b) and Wimbauer et al. (1997a), the spatial organisation of ON and OFF LGN inputs to the model simple cell is established by a competitive process acting over small spatial offsets and facilitation over larger offsets. In contrast, all input populations in our model facilitate the development of connectivity patterns from other input populations, with spatial offsets determined by a temporal delay. The implication this has for the development of ON and OFF inputs is that their segregation in space depends on the distribution of relative response latencies between them, and thus the distribution of luminance contours and motion velocities in the visual environment. Our model therefore allows for the possibility that not all simple cells have well defined ON and OFF subfields, and that these inputs may overlap substantially in some cells. While there is some evidence of overlapping ON and OFF inputs in V1 (Mata & Ringach, 2005), it is difficult to tell whether this is a result of converging feedforward input, interactions within the

cortex, or even a difficulty in distinguishing simple cells from other cell classes, such as complex cells (reviewed in Carandini et al., 2005). Despite these confounding factors, controlling the visual environment during development may at least modify the distribution of ON-OFF overlap over a large population of simple cells. Such experiments may help to verify the applicability of our model to the development of ON and OFF inputs to simple cells.

The fact that orientation selectivity is established, to some degree, at the time of eye opening (Chapman & Stryker, 1993; Krug et al., 2001; White et al., 2001; Li et al., 2006), suggests that ON and OFF inputs are already partially organised at that time, and thus that cortical RFs begin to refine before the onset of normal visual experience. The model of Miller (1994) utilises the centre-surround structure of LGN RFs to yield such cortical RF structures, which is possible even activity in the inputs is not patterned. However, Oshiro & Weliky (2006) demonstrated that the RFs of neurons in the LGN lack this centre-surround structure in the early stages of visual development. But in several animals, patterned activity is present before the onset of vision in the form of retinal waves, which are active up until the time of eye opening (Huberman et al., 2008b). During the final stage of retinal waves, ON and OFF retinal ganglion cells participate in waves with a time delay of approximately one second (Kerschensteiner & Wong, 2008). It is very possible that these delays are recapitulated in the developing LGN. In Chapter 5, we proposed that STDP rules with long time constants may be appropriate for the interaction with retinal waves and play a role in the refinement of RFs in the superior colliculus. If long time scale STDP rules were also at play at geniculocortical synapses, then an asymmetric STDP rule with a decay time constant of 0.3s would provide a phase shift of 0.35 cycles, and a symmetric rule would provide a shift of 0.45 cycles. Thus, the ON-OFF delay during late stage retinal waves would, in principal, be suitable for instructing the segregation of ON and OFF inputs to V1.

Extensions to our model

Our model, as it stands, provides an outline for a novel way in which STDP and motion stimuli can establish direction selectivity in geniculorecipient cells. Our analytical results make predictions regarding the development of simple cells, vis-à-vis the dependence of spatial frequency tuning, and independence of spatiotemporal tilt, in STRFs on the stimulus velocity. However, an exposition of the model that more faithfully captures the biological details of the geniculocortical circuit is required before its predictive power for simple cell development can be ascertained. To do this, the following extensions to the model would be necessary. First, the spatial and temporal response profile of LGN cells to visual stimuli, though immature, have distinct waveforms at the time of eye opening (Tavazoie & Reid, 2000) and will enter our derivation for the synaptic strength dynamics through $\alpha(t)$, which will become $\alpha(x, t)$. The band pass nature of LGN spatiotemporal RFs are likely to constrain the range of spatial frequencies that may develop in the connectivity patterns. Second, we must extend our input layer into two dimensions, and constrain the connectivity pattern to within a spatially restricted arbour, so as to confirm whether realistic simple cell RFs will develop when training stimuli are presented with motion in multiple directions. Finally, the observation that all orientations and directions are encoded preferentially over the V1 population (Hubel & Wiesel, 1968; Daw & Wyatt, 1976) must be accounted for in our model, requiring that multiple output neurons acquire unique STRFs, given the same set of training stimuli. This may be achieved by seeding the initial output RF with some randomness. Additional non-linearities, either in the response characteristics of the output cell, or in the plasticity rule, may further help to break any symmetries in the training stimuli and encourage different STRFs to develop.

It is reasonable to think that the framework we have presented here may be ap-

plied to other units of computation that have been proposed to comprise the visual pathway. For example, alternative models for DS simple cells posit that inputs come from multiple non-DS simple cells, with RFs that cover nearby retinotopic locations or that have phase shifted ON and OFF subfields, and which are delayed with respect to each other. Models of complex cells, on the other hand, may arise from the convergence of DS simple cells. Incorporating these cell types into the input layer of our framework would be a simple matter of appropriately defining $\alpha(x, t)$, so as to capture the spatiotemporal response properties of non-DS simple cells, on the one hand, or DS simple cells, on the other.

Conclusions

Summary of aims and results

Travelling wave patterns of neuronal activity are a prominent feature in the developing brain, resulting from spontaneous depolarisation in highly excitable circuits (Feller, 1999; Ben-Ari, 2001; Blankenship & Feller, 2010), or in response to sensory stimuli (Frechette et al., 2005). The strong inter-neuronal correlations induced by travelling waves have long been proposed to act through Hebbian mechanisms of synaptic plasticity (Hebb, 1949) and drive changes in the synaptic connectivity of neural networks that underly their refinement during early development. Early theoretical formulations of Hebbian plasticity, based on time averaged firing rates, have demonstrated the influence of spatially structured correlations on plastic networks (Kohonen, 1974; Hopfield, 1982; Linsker, 1986a-b; Goodhill, 1993; Miller & MacKay, 1994). However, the spatial structure of correlations induced by travelling waves are tightly coupled with the temporal structure of correlations, which can vary over time scales as short as tens of milliseconds. Spike-timing dependent plasticity (STDP) is an expression of Hebbian plasticity that is sensitive to correlations over these time scales (Markram et al., 1997; Zhang et al., 1998; Kempter et al., 1999; Song et al., 2000), but despite its application to several specific paradigms involving travelling waves (Buchs & Senn, 2002; Rao & Sejnowski, 2003; Wenisch et al., 2005; Yu et al., 2006; Honda et al., 2011), a more general and thorough investigation of the interaction between STDP and travelling waves, and its influence on developing neural networks, has been lacking.

Here, we have examined this interaction in the framework of a simple yet biologically plausible feedforward network of spiking neurons. Our results, derived analytically and tested in simulations, describe a novel and robust pattern forming process by which STDP can build structure into feedforward circuits, and fall into a wider class of pattern formation that has had broad application to other areas of biology (Mein-

hardt & Giesler, 2000; Kondo & Miura, 2010). Moreover, we have identified distinct roles for properties such as the wave speed, wave size, and the decay time of STDP rules, in the pattern forming process, which we have linked directly with the refinement of postsynaptic receptive fields (RFs). Our model generalises across different brain areas, and we have applied our results to two specific cases of development in the early visual system, including retinotopic refinement in the superior colliculus (SC) during spontaneously active retinal waves, and the development of direction selective simple cells in primary visual cortex (V1) during early visual experience, yielding predictions that are testable using existing experimental paradigms.

Pattern formation with spike-timing dependent plasticity

The pattern formation process described in Chapter 3 is analogous to a Turing instability in reaction-diffusion systems (Turing, 1952), whereby an initial homogenous distribution of synaptic strengths becomes unstable and the fastest growing eigenmode dominates, resulting in a periodic connectivity pattern. We identified two key properties of circuit dynamics that influence this pattern: the speed of travelling waves, and the decay times that characterise the STDP rule. Increases to either of these decrease the spatial frequency of the pattern, and thus link our theoretical prediction to experimental observables.

One useful assumption in our derivation posits that pairs of input and output spikes contribute linearly to changes in synaptic strength (Kempster et al., 1999, Song et al., 2000). However, STDP is influenced by a number of other factors that can be modulated by activity patterns with multiple spikes. Several models have stepped beyond the phenomenology of STDP and addressed biophysical processes responsible for it. These models may not necessarily have the same interaction with travelling waves, as we found in Chapter 5 when modelling the STDP_{sub} rule, a proxy for the calcium

dependent model in Graupner & Brunel (2012). Nevertheless, a great deal remains to be worked out regarding the spike-timing dependence of plasticity (Feldman, 2012), and synaptic changes that result from complex wave patterns, which recruit multiple synaptic inputs, is unknown. Our results provide a tractable platform from which to explore this interaction in greater detail. A major prediction from our work is that, for the process of pattern formation to serve the developing brain, synaptic plasticity must favour synaptic weakening, to induce competition, and must involve mechanisms of weakening that operate over longer time scales than mechanisms of strengthening.

A role for pattern formation in the development of receptive fields

An important application of pattern formation is the development of sensory RFs. In particular, in Chapter 4, we demonstrated how a RF can be likened to a single cycle in the periodic patterns that result from the interaction of travelling waves with STDP. In our model, the wave speed controls the spatial frequency of the periodic pattern, and therefore the developmental trajectory of the RF, whereby increasing the wave speed increases the mature size of the RF or its component subfields. This motivates one way in which our theoretical predictions can be tested experimentally.

In ferrets and mice, early visual experience is required for the proper development of neurons in the lateral geniculate nucleus (LGN). In dark reared ferrets, LGN RFs are more diffuse and multi-peaked (Akerman et al., 2004), and in dark reared mice, synaptic inputs from the retina become more numerous after an initial period of pruning (Hooks & Chen, 2006; Hooks & Chen). Similarly, in the hamster SC, dark rearing leads to a loss of refinement in the form of enlarged RFs (Carrasco et al., 2005). These results suggest critical periods of development during which the refine-

ment, or maintenance, of RFs in the LGN and SC are sensitive to the provision of visual experience. What is not yet known is whether vision is permissive or instructive in these areas. This could be tested by training these animals on visual stimuli with moving contours, such as high contrast bars, during these critical periods. If motion is along just one axis, our model predicts that RFs would become broader or narrower, along that axis only, with faster or slower stimulus speeds, respectively. This prediction also depends on the type of plasticity rule operating at these synapses, which is an area of current investigation. Another suitable model animal is *Xenopus Laevis*, for which asymmetric STDP rules at retinotectal synapses have been recorded (Zhang et al., 1998). RFs in the tectum of *Xenopus* tadpoles are also highly plastic. When presented with visual stimuli consisting of bars moving in a single direction, RF centres in the tectum shift in the opposite direction (Richards et al., 2010), consistent with our simulations using an asymmetric STDP rule in Chapter 3.

Potential roles for retinal waves in retinocollicular refinement

In Chapter 5, we found that the type of plasticity rule determines which features of retinal waves are important for the refinement of RFs. We constrained our investigation to a mode of synaptic plasticity that is coincident with retinal waves, burst-timing dependent plasticity (BTDP) (Butts et al., 2007), and proposed two STDP rules that might give rise to BTDP. One of these was the symmetric STDP rule used throughout this thesis (Eq. 2.4 in General Methods). This rule required wave propagation over sufficiently large distances to refine RFs to a smaller size, and expanded the RF if waves travelled distances shorter than this. The STDP_{sub} rule (see Methods in Chapter 5), by comparison, had the opposite effect. Waves that propagated large distances recruited more inputs per unit time, which resulted in RF growth, whereas waves that propagated short distances recruited fewer inputs,

resulting in refinement of the RF to a smaller size.

A major challenge in determining which retinal wave properties are important for development is in reliably controlling those properties *in vivo*. Two recent studies using mice, one in which wave propagation was stunted (Xu et al., 2011), and another in which pan-retinal stimulation was applied repetitively (Zhang et al., 2012), both produced similar levels of retinocollicular refinement observed in control mice, despite the very different structure of correlations that would be expected from these types of activity. Because of the geometry of the retina and translucent optics of the eye at this stage of development, the measurement of activity in the retina *in vivo* poses another considerable challenge in deciphering a role for retinal waves in refinement. Whether the expected patterns of perturbed retinal activity were present *in vivo* in these particular studies is unknown.

To date, *in vitro* experimental recordings of unperturbed retinal waves have been made in several different animals by different teams (Feller et al., 1996; Sernagor et al., 2003; 2000sernagor.odonovan; Stafford et al., 2009; Maccione et al., 2014). These exhibit statistical properties, such as wave frequency and speed, that are shared with *in vivo* recordings of patterned activity downstream of the retina (Ackman et al., 2012), adding weight to the likelihood that *in vitro* retinal waves are representative of those *in vivo*. A measure of the correlations induced by these waves is key to understanding how they may instruct refinement through Hebbian plasticity. Previous analyses of retinal waves had focussed only on spatial correlations, which were found to decay as a function of distance between RGCs (Wong et al., 1993; Stafford et al., 2009). However, to elucidate whether unperturbed retinal waves can instruct refinement via pattern formation, we had to analyse both spatial and temporal correlations. We therefore developed a new method of analyses using the sCCG, and found that simulated stage 2 retinal waves cast space-time inseparable correlations

that could instruct RF refinement in the manner outlined in Chapter 4. We have also conducted a preliminary analysis of experimentally recorded stage 2 retinal waves, using the sCCG, and found a correlation structure similar to those of simulated waves. We intend to extend this analysis to a much larger set of data that records changing activity properties throughout stage 2 retinal waves (Maccione et al., 2014). This approach will help to determine which retinal wave features may be instructive for refinement in the SC, according to different models of Hebbian plasticity, with the aim to guide and interpret future experiments in this field.

Development of selectivity in simple cells

Our theoretical results in Chapter 6 offer a new framework for investigating the development of orientation and direction selectivity in V1 simple cells. In particular, it differs from previous models of simple cell development in two ways. Instead of ON and OFF inputs from the LGN competing for visuotopic representation (Linsker, 1986b; Miller, 1994; Ohshiro & Weliky, 2006), we show how they can mutually facilitate growth at separate locations due to their different response latencies to moving visual stimuli. We predict that the segregation of ON and OFF subfields would increase linearly with this temporal offset. A second strong prediction is that the development of direction selectivity, while requiring exposure to visual motion, should not in general depend on the speed or directional bias of the training stimuli, features that were problematic in previous models (Wimbauer et al., 1997a; Blais et al., 2000). This second prediction is a consequence of a previously discussed feature of our model: the spatial frequency of the simple cell RF results from periodic pattern formation in our model, and therefore does depend on the speed of the training stimulus.

These three predictions can be tested using experimental techniques that are already

well established, including the presentation of training stimuli such as high contrast moving bars, and the mapping of matured RFs by reverse correlation with noise stimuli (DeAngelis et al., 1993a; Li et al., 2008; van Hooser et al., 2012). The precise details of these predictions, however, will require the addition of several details to our model, including the spatiotemporal stimulus response properties of geniculate neurons, constraints on the feedforward anatomy in two dimensions, and simulations with more complex training stimuli. This work is already in progress.

References

- Abbott, L. F., Nelson, S. B. (2000), Synaptic plasticity: taming the beast, *Nat. Neurosci. Suppl.* 3, 1178-1183
- Ackman, J. B., Burbridge, T. J., Crair, M. C. (2012), Retinal waves coordinate patterned activity throughout the developing visual system, *Nature* 490, 219-225
- Adelson, E. H., Bergen, J. R. (1985), Spatiotemporal energy models for the perception of motion, *J. Opt. Soc. Am. A* 2, 284-299
- Akerman, C. J., Grubb, M. S., Thompson, I. D. (2004), Spatial and Temporal Properties of Visual Responses in the Thalamus of the Developing Ferret, *J. Neurosci.* 24, 170-182
- Akerman, C. J., Smyth, D., Thompson, I. D. (2002), Visual Experience before Eye-Opening and the Development of the Retinogeniculate Pathway, *Neuron* 36, 869-879
- Albrecht, D. G., Geisler, W. S. (1991), Motion selectivity and the contrast-response function of simple cells in the visual cortex, *Vis. Neurosci.* 7, 531-546
- Alonso, J. M., Usrey, W. M., Reid, R. C. (2001), Rules of Connectivity between Geniculate Cells and Simple Cells in Cat Primary Visual Cortex, *J. Neurosci.* 21, 4002-4015
- Angelucci, A., Levitt, J. B., Lund, J. S. (2002), Anatomical origins of the classical receptive field and modulatory surround field of single neurons in macaque visual cortical area V1,

Baker, P. M., Bair, W. (2012), Inter-neuronal correlation distinguishes mechanisms of direction selectivity in cortical circuit models, *J. Neurosci.* 32, 8800-8816

Bansal, A., Singer, J. H., Hwang, B. J., Xu, W., Beaudet, A., Feller, M. B. (2002), Mice Lacking Specific Nicotinic Acetylcholine Receptor Subunits Exhibit Dramatically Altered Spontaneous Activity Patterns and Reveal a Limited Role for Retinal Waves in Forming ON and OFF Circuits in the Inner Retina, *J. Neurosci.* 20, 7672-7681

Barlow, H. B., Levick, W. R. (1965), The mechanism of directionally selective units in rabbit's retina, *J. Physiol.* 178, 477-504

Ben-Ari, Y. (2001), Developing networks play a similar melody, *Trends Neurosci.* 24, 353-360

Bender, V. A., Bender, K. J., Brasier, D. J., Feldman, D. E. (2006), Two Coincidence Detectors for Spike Timing-Dependent Plasticity in Somatosensory Cortex, *J. Neurosci.* 26, 4166-4177

Bi, G. Q., Poo, M. M. (1998), Synaptic Modifications in Cultured Hippocampal Neurons: Dependence on Spike Timing, Synaptic Strength, and Postsynaptic Cell Type, *J. Neurosci.* 18, 10464-10472

Bienenstock, E. L., Cooper, L. N., Munro, P. W. (1982), Theory for the development of neuron selectivity: orientation specificity and binocular interaction in visual cortex, *J. Neurosci.* 2, 32-48

Blais, B., Cooper, L. N., Shouval, H. (2000), Formation of Direction Selectivity in Natural Scene Environments, *Neural Computation* 12, 1057-1066

Blakemore, C., Cooper, G. F. (1970), Development of the Brain depends on the Visual Environment, *Nature* 228, 477-478

Blakemore, C., Tobin, E. (1972), Lateral Inhibition Between Orientation Detectors in the Cat's Visual Cortex, *Exp. Brain Res.* 15, 439-440

- Blakemore, C., Mitchell, D. E. (1973), Environmental Modification of the Visual Cortex and the Neural Basis of Learning and Memory, *Nature* 241, 467-468
- Blakemore, C., Vital-Durand, F. (1986a), Organization and post-natal development of the monkey's lateral geniculate nucleus, *J. Physiol.* 380, 453-491
- Blakemore, C., Vital-Durand, F. (1986b), Effects of visual deprivation on the development of the monkey's lateral geniculate nucleus, *J. Physiol.* 380, 493-511
- Blakemore, C., van Sluyters, R. C. (1975), Innate and environmental factors in the development of the kitten's visual cortex, *J. Physiol.* 248, 663-716
- Blankenship, A. G., Feller, M. B. (2010), Mechanisms underlying spontaneous patterned activity in developing neural circuits, *Nat. Rev. Neurosci.* 11, 18-29
- Bodenstein, L., Sidman, R. L. (1987), Growth and Development of the Mouse Retinal Pigment Epithelium, *Develop. Biol.* 121, 192-204
- Braastad, B. O., Heggelund, P. (1985), Development of Spatial Receptive-Field Organization and Orientation Selectivity in Kitten Striate Cortex, *J. Neurophysiol.* 53, 1158-1178
- Buchs, N. J., Senn, W. (2002), Spike-Based Synaptic Plasticity and the Emergence of Direction Selective Simple Cells: Simulation Results, *J. Comp. Neurosci.* 13, 167-186
- Butts, D. A., Feller, M. B., Shatz, C. J., Rokhsar, D. S. (1999), Retinal waves are governed by collective network properties, *J. Neurosci.*, 19, 3580-3593
- Butts, D. A., Kanold, P. O., Shatz, C. J. (2007), A Burst-Based "Hebbian" Learning Rule at Retinogeniculate Synapses Links Retinal Waves to Activity-Dependent Refinement, *Plos Biol.* 5, e61
- Butts, D. A., Rokhsar, D. S. (2001), The information content of spontaneous retinal waves, *J. Neurosci.* 21, 961-973
- Cai, D., DeAngelis, G. C., Freeman, R. D. (1997), Spatiotemporal Receptive Field Organization in the Lateral Geniculate Nucleus of Cats and Kittens, *J. Neurophysiol.* 78,

Cai, Y., Gavornik, J. P., Cooper, L. N., Yeung, L. C., Shouval, H. Z. (2006), Effect of Stochastic Synaptic and Dendritic Dynamics on Synaptic Plasticity in Visual Cortex and Hippocampus, *J. Neurophysiol.* 97, 375-386

Cang, J., Rentería, R. C., Kaneko, M., Liu, X., Copenhagen, D. R., Stryker, M. P. (2005), Development of Precise Maps in Visual Cortex Requires Patterned Spontaneous Activity in the Retina, *Neuron* 48, 797-809

Cang, J., Niell, C. M., Liu, X., Pfeifferberger, C., Feldheim, D. A., Stryker, M. P. (2008a), Selective disruption of one Cartesian axis of cortical maps and receptive fields by deficiencies in ephrin-As and structured activity, *Neuron* 57, 511-523

Cang, J., Wang, L., Stryker, M. P., Feldheim, D. A. (2008b), Roles of ephrin-As and structured activity in the development of functional maps in the superior colliculus, *J. Neurosci.* 28, 11015-11023

Cang, J., Feldheim, D. A. (2013), Developmental Mechanisms of Topographic Map Formation and Alignment, *Annu. Rev. Neurosci.* 36, 51-77

Caporale, N., Dan, Y. (2008), Spike timing-dependent plasticity: a Hebbian learning rule, *Annu. Rev. Neurosci.* 31, 25-46

Carandini, M., Demb, J. B., Mante, V., Tolhurst, D. J., Dan, Y., Olshausen, B. A., Gallant, J. L., Rust, N. C. (2005), Do we know what the early visual system does?, *J. Neurosci.* 25, 10577-10597

Carrasco, M. M., Razak, K. A., Pallas, S. L. (2005), Visual Experience Is Necessary for Maintenance But Not Development of Receptive Fields in Superior Colliculus, *J. Neurophysiol.* 94, 1962-1970

Chalupa, L. M. (2009), Retinal waves are unlikely to instruct the formation of eye-specific retinogeniculate projections, *Neural Development* 4, 25

Chance, F. S., Nelson, S. B., Abbott, L. F. (1998), Synaptic Depression and the Temporal Response Characteristics of V1 Cells, *J. Neurosci.* 18, 4785-4799

- Chandrasekaran, A. R., Plas, D. T., Gonzalez, E., Crair, M. C. (2005), Evidence for an Instructive Role of Retinal Activity in Retinotopic Map Refinement in the Superior Colliculus of the Mouse, *J. Neurosci.* 25, 6929-6938
- Chandrasekaran, A. R., Furuta, Y., Crair, M. C. (2009), Consequences of axon guidance defects on the development of retinotopic receptive fields in the mouse colliculus, *J. Physiol.* 587, 953-963
- Chapman, B., Stryker, M. P. (1993), Development of orientation selectivity in ferret visual cortex and effects of deprivation, *J. Neurosci.* 13, 5251-5256
- Chen, C., Regehr, W. G. (2000), Developmental Remodeling of the Retinogeniculate Synapse, *Neuron* 28, 955-966
- Chklovskii, D. B., Mel, B. W., Svoboda, K. (2004), Cortical wiring and information storage, *Nature* 431, 782-788
- Clopath, C., Büsing, L., Vasilaki, E., Gerstner, W. (2010), Connectivity reflects coding: a model of voltage-based STDP with homeostasis, *Nat. Neurosci.* 13, 344-352
- Conhaim, J., Cedarbaum, E. R., Barahimi, M., Moore, J. G., Becker, M. I., Gleiss, H., Kohl, C., Moody, W. J. (2010), Bimodal Septal and Cortical Triggering and Complex Propagation Patterns of Spontaneous Waves of Activity in the Developing Mouse Cerebral Cortex, *Develop. Neurobiol.*, 70, 679-692
- Cynader, M., Berman, N., Hein, A. (1975), Cats Raised in a One-Directional World: Effects on Receptive Fields in Visual Cortex and Superior Colliculus, *Exp. Brain Res.* 22, 267-280
- Daniels, J. D., Pettigrew, J. D., Norman, J. L. (1978), Development of single-neuron responses in kitten's lateral geniculate nucleus, *J. Neurophysiol.* 41, 1373-1393
- Daw, N. W., Wyatt, H. J. (1976), Kittens reared in a unidirectional environment: Evidence for a critical period, *J. Physiol.* 257, 155-170
- De Valois, R. L., Cottaris, N. P., Mahon, L. E., Elfar, S. D., Wilson, J. A. (2000), Spatial and temporal receptive fields of geniculate and cortical cells and directional selectivity,

- DeAngelis, G. C., Freeman, R. D., Ohzawa, I. (1994), Length and Width Tuning of Neurons in the Cat's Primary Visual Cortex, *J. Neurophysiol.* 71, 347-374
- DeAngelis, G. C., Ohzawa, I., Freeman, R. D. (1993a), Spatiotemporal organization of simple-cell receptive fields in the cat's striate cortex. I. General characteristics and postnatal development, *J. Neurophysiol.* 69, 1091-1117
- DeAngelis, G. C., Ohzawa, I., Freeman, R. D. (1993b), Spatiotemporal organization of simple-cell receptive fields in the cat's striate cortex. II. Linearity of Temporal and Spatial Summation, *J. Neurophysiol.* 69, 1118-1135
- Delgado, J. Y., Gómez-González, J. F., Desai, N. S. (2010), Pyramidal neuron conductance state gates spike-timing-dependent plasticity, *J. Neurosci.* 30, 15713-15725
- Dhande, O. S., Hua, E. W., Guh, E., Yeh, J., Bhatt, S., Zhang, Y., Ruthazer, E. S., Feller, M. B., Crair, M. C. (2011), Development of single retinofugal axon arbours in normal and $\beta 2$ knock-out mice, *J. Neurosci.* 31, 3384-3399
- Dong, D. W., Atick, J. J. (1995), Statistics of natural time-varying images, *Network: Comput. Neural Syst.* 6, 345-358
- Dunn, T. A., Wang, C. T., Colicos, M. A., Zaccolo, M., DiPilato, L. M., Zhang, J., Tsien, R. Y., Feller, M. B. (2006), Imaging of cAMP Levels and Protein Kinase A Activity Reveals That Retinal Waves Drive Oscillations in Second-Messenger Cascades, *J. Neurosci.* 26, 12807-12815
- Espinoso, J. S., Stryker, M. P. (2012), Development and Plasticity of the Primary Visual Cortex, *Neuron* 75, 230-249
- Etherington, S. J., Williams, S. R. (2011), Postnatal development of intrinsic and synaptic properties transforms signals in the layer 5 excitatory neural network of the visual cortex, *J. Neurosci.* 31, 9526-9537
- Fagiolini, M., Pizzorusso, T., Berardi, N., Domenici, L., Maffei, L. (1994), Functional Postnatal Development of the Rat Primary Visual Cortex and the Role of Visual Experience:

Dark Rearing and Monocular Deprivation, *Vision Res.* 34, 709-720

Feidler, J. C., Saul, A. B., Murthy, A., Humphrey, A. L. (1997), Hebbian learning and the development of direction selectivity: the role of geniculate response timings, *Network: Comput. Neural Syst.* 8, 195-214

Feldman, D. E. (2000), Timing-based LTP and LTD at vertical inputs to layer II/III pyramidal cells in rat barrel cortex, *Neuron* 27, 45-56

Feldman, D. E. (2012), The spike-timing dependence of plasticity, *Neuron* 75, 556-571

Feller, M. (1999), Spontaneous Correlated Activity in Developing Neural Circuits, *Neuron* 22, 653-656

Feller, M. B. (2009), Retinal waves are likely to instruct the formation of eye-specific retinogeniculate projections, *Neural Development* 4, 24

Feller, M. B., Butts, D. A., Aaron, H. L., Rokhsar, D. S., Shatz, C. J. (1997), Dynamics processes shape spatiotemporal properties of retinal waves, *Neuron* 19, 293-306

Feller, M. B., Wellis, D. P., Stellwagen, D., Werblin, F. S., Shatz, C. J. (1996), Requirement for cholinergic synaptic transmission in the propagation of spontaneous retinal waves, *Science* 272, 1182-1187

Fiete, I. R., Senn, W., Wang, C. Z. H., Hahnloser, R. H. R. (2010), Spike-time-dependent plasticity and heterosynaptic competition organize networks to produce long scale-free sequences of neural activity, *Neuron* 65, 563-576

Ford, K. J., Félix, A. L., Feller, M. B. (2012), Cellular mechanisms underlying spatiotemporal features of cholinergic retinal waves, *J. Neurosci.* 32, 850-863

Fortin, S., Chabli, A., Dumont, I., Shumikhina, S., Itaya, S. K., Molotchnikoff, S. (1999), Maturation of visual receptive field properties in the rat superior colliculus, *Dev. Brain Res.* 112, 55-64

Frechette, E. S., Sher, A., Grivich, M. I., Petrusca, D., Litke, A. M., Chichilnisky, E. J. (2005), Fidelity of the Ensemble Code for Visual Motion in Primate Retina, *J. Neuro-*

physiol. 94, 119-135

Froemke, R. C., Dan, Y. (2002), Spike-timing-dependent synaptic modification induced by natural spike trains, *Nature* 416, 433-438

Froemke, R. C., Poo, M. M., Dan, Y. (2005), Spike-timing-dependent synaptic plasticity depends on dendritic location, *Nature* 434, 221-225

Froemke, R. C., Tsay, I. A., Raad, M., Long, J. D., Dan, Y. (2006), Contribution of Individual Spikes in Burst-Induced Long-Term Synaptic Modification, *J. Neurophysiol.* 95, 1620-1629

Fuenzalida, M., de Sevilla, D. F., Buño, W. (2007), Changes of the EPSP waveform regulate the temporal window for spike-timing-dependent plasticity, *J. Neurosci.* 27, 11940-11948

Gütig, R., Aharonov, R., Rotter, S., Sompolinsky, H. (2003), Learning Input Correlations through Nonlinear Temporally Asymmetric Hebbian Plasticity, *J. Neurosci.* 23, 3697-3714

Garaschuk, O., Linn, J., Eilers, J., Konnerth, A. (2000), Large-scale oscillatory calcium waves in the immature cortex, *Nature* 3, 452-459

Gerstner, W., Kempter, R., van Hemmen, J. L., Wagner, H. (1996), A neuronal learning rule for sub-millisecond temporal coding, *Nature* 383, 76-78

Gjorgjieva, J., Clopath, C., Audet, J., Pfister, J. P., (2011), A triplet spike-timing-dependent plasticity model generalizes the Bienenstock-Cooper-Munro rule to higher-order spatiotemporal correlations, *Proc. Natl. Acad. Sci. USA* 108, 19383-19388

Gjorgjieva, J., Toyoizumi, T., Eglen, S. J. (2009), Burst-time-dependent plasticity robustly guides ON/OFF segregation in the lateral geniculate nucleus, *PLoS Comp. Biol.* 5(12): e1000618

Goodhill, G. J. (1993), Topography and ocular dominance: a model exploring positive correlations, *Biol. Cybern.* 69, 109-118

- Graupner, M., Brunel, N. (2012), Calcium-based plasticity model explains sensitivity of synaptic changes to spike pattern, rate, and dendritic location, *Proc. Natl. Acad. Sci. USA* 109, 3991-3996
- Grubb, M. S., Rossi, F. M., Changeux, J. P., Thompson, I. D. (2003), Abnormal functional organization in the dorsal lateral geniculate nucleus of mice lacking the $\beta 2$ subunit of the nicotinic acetylcholine receptor, *Neuron* 40, 1161-1172
- Guo, Y., Huang, S., de Pasquale, R., McGehrin, K., Lee, H. K., Zhao, K., Kirkwood, A. (2012), Dark Exposure Extends the Integration Window for Spike-Timing-Dependent Plasticity, *J. Neurosci.* 32, 15027-15035
- Gur, M., Kagan, I., Snodderly, D. M. (2005), Orientation and Direction Selectivity of Neurons in V1 of Alert Monkeys: Functional Relationships and Laminar Distributions, *Cerebral Cortex* 15, 1207-1221
- Hao, J., Oertner, T. G. (2012), Depolarization gates spine calcium transients and spike-timing-dependent potentiation, *Current Opinions in Neurobiology* 22, 509-515
- Hebb, D. O. (1949), The organization of Behaviour, Wiley, New York
- Heifets, B. D., Castillo, P. E. (2009), Endocannabinoid Signaling and Long-Term Synaptic Plasticity, *Annu. Rev. Physiol.* 71, 283-306
- Henderson, Z., Finlay, B. L., Wikler, K. C. (1986), Development of Ganglion Cell Topography in Ferret Retina, *J. Neurosci.* 8, 1194-1205
- Honda, H., Urakubo, H., Tanaka, K., Kuroda, S. (2011), Analysis of development of direction selectivity in retinotectum by a neural circuit model with spike timing-dependent plasticity, *J. Neurosci.* 31, 1516-1527
- Hooks, B. M., Chen, C. (2006), Distinct Roles for Spontaneous and Visual Activity in Remodeling of the Retinogeniculate Synapse, *Neuron* 52, 281-291
- Hooks, B. M., Chen, C. (2008), Vision Triggers an Experience-Dependent Sensitive Period at the Retinogeniculate Synapse, *J. Neurosci.* 28, 4807-4817

- Hopfield, J. J. (1982), Neural networks and physical systems with emergent collective computational abilities, *Proc. Natl. Acad. Sci. USA* 79, 2554-2558
- Hubel, D. H., Wiesel, T. N. (1959), Receptive fields of single neurons in the cat's striate cortex, *J. Physiol.* 148, 574-591
- Hubel, D. H., Wiesel, T. N. (1962), Receptive fields, binocular interaction and functional architecture in the cat's visual cortex, *J. Physiol.* 160, 106-154
- Hubel, D. H., Wiesel, T. N. (1963), Receptive fields of cells in striate cortex of very young, visually inexperienced kittens, *J. Neurophysiol.* 26, 994-1002
- Hubel, D. H., Wiesel, T. N. (1968), Receptive fields and functional architecture of monkey striate cortex, *J. Physiol.* 195, 215-243
- Huberman, A. D., Feller, M. B., Chapman, B. (2008b), Mechanisms Underlying Development of Visual Maps and Receptive Fields, *Annu. Rev. Neurosci.* 31, 479-509
- Huberman, A. D., Manu, M., Koch, S. M., Susman, M. W., Lutz, A. B., Ullian, E. M., Baccus, S. A., Barres, B. A. (2008a), Architecture and Activity-Mediated Refinement of Axonal Projections from a Mosaic of Genetically Identified Retinal Ganglion Cells, *Neuron* 59, 425-438
- Humphrey, A. L., Saul, A. B. (1998), Strobe Rearing Reduces Direction Selectivity in Area 17 by Altering Spatiotemporal Receptive-Field Structure, *J. Neurophysiol.* 80, 2991-3004
- Jester, J. V. (2008), Corneal crystallins and the development of cellular transparency, *Seminars in Cell and Developmental Biology* 19, 82-93
- Jones, J. P., Palmer, L. A. (1987), An Evaluation of the Two-Dimensional Gabor Filter Model of Simple Receptive Fields in Cat Striate Cortex, *J. Neurophysiol.* 58, 1233-1258
- Karmarkar, U. R., Buonomano, D. V. (2002), A Model of Spike-Timing Dependent Plasticity: One or Two Coincidence Detectors?, *J. Neurophysiol.* 88, 507-513

- Katz, L. C., Shatz, C. J. (1996), Synaptic activity and the construction of cortical circuits, *Science* 274, 1133-1138
- Kempler, R., Gerstner, W., van Hemmen, J. L. (1999), Hebbian learning and spiking neurons, *Phys. Rev. E* 59, 4498-4514
- Kerschensteiner, D. Wong, R. O. L. (2008), A precisely timed asynchronous pattern of ON and OFF retinal ganglion cell activity during propagation of retinal waves, *Neuron* 58, 851-858
- Kohonen, T. (1974), An Adaptive Associative Memory Principle, *IEEE Trans. on Computers* C-23, 444-445
- Kolodkin, A. L., Tessier-Lavigne, M. (2011), Mechanisms and Molecules of Neuronal Wiring: A Primer, *Cold Spring Harb. Perspect. Biol.* 3, a001727
- Kondo, S., Miura, T. (2010), Reaction-Diffusion Model as a Framework for Understanding Biological Pattern Formation, *Science* 329, 1616-1620
- Krug, K., Akerman, C. J., Thompson, I. D. (2001), Responses of neurons in neonatal cortex and thalamus to patterned visual stimulation through the naturally closed lids, *J. Neurophysiol.* 85, 1436-1443
- Leinekugel, X., Khalilov, I., Ben-Ari, Y., Khazipov, R. (1998), Giant depolarizing potentials: the septal pole of the hippocampus paces the activity of the developing intact septohippocampal complex in vitro, *J. Neurosci.* 18, 6349-6357
- Li, Y., van Hooser, S. D., Mazurek, M., White, L. E., Fitzpatrick, D. (2008), Experience with moving visual stimuli drives the early development of cortical direction selectivity, *Nature* 456, 952-956
- Li, Y., Fitzpatrick, D., White, L. E. (2006), The development of direction selectivity in ferret visual cortex requires early visual experience, *Nat. Neurosci.* 9, 676-681
- Linsker, R. (1986a), From basic network principles to neural architecture: Emergence of spatial-opponent cells, *Proc. Natl. Acad. Sci. USA* 83, 7508-7512

- Linsker, R. (1986b), From basic network principles to neural architecture: Emergence of orientation-selective cells, *Proc. Natl. Acad. Sci. USA* 83, 8390-8394
- Liu, X. B., Murray, K. D., Jones, E. G. (2004), Switching of NMDA Receptor 2A and 2B Subunits at Thalamic and Cortical Synapses during Early Postnatal Development, *J. Neurosci.* 24, 8885-8895
- Liu, X., Chen, C. (2008), Different Roles for AMPA and NMDA Receptors in Transmission at the Immature Retinogeniculate Synapse, *J. Neurophysiol.* 99, 629-643
- Maccione, A., Hennig, M. H., Gandolfo, M., Muthmann, O., van Coppenhagen, J., Eglén, S. J., Berdondini, L., Sernagor E. (2014), Following the ontogeny of retinal waves: pan-retinal recordings of population dynamics in the neonatal mouse, *J. Physiol.* 592, 1545-1563
- Markram, H., Gerstner, W., Sjöström, P. J. (2011), A history of spike-timing-dependent plasticity, *Front. Synaptic Neurosci.* 3, 4
- Markram, H., Lübke, J., Frotscher, M., Sakmann, B. (1997), Regulation of Synaptic Efficacy by Coincidence of Postsynaptic APs and EPSPs, *Science* 275, 213-215
- Mata, M. L., Ringach, D. L. (2005), Spatial Overlap of ON and OFF Subregions and Its Relation to Response Modulation Ratio in Macaque Primary Visual Cortex, *J. Neurophysiol.* 93, 919-928
- McCormick, D. A., Prince, D. A. (1987), Post-natal development of electrophysiological properties of rat cerebral cortical pyramidal neurones, *J. Physiol.* 393, 743-762
- McLaughlin, T., Torborg, C. L., Feller, M. B., O’Leary, D. D. M. (2003a), Retinotopic map refinement requires spontaneous retinal waves during a brief critical period of development, *Neuron* 40, 1147-1160
- McLaughlin, T., Hindges, R., Yates P. A., O’Leary, D. D. M. (2003b), Bifunctional action of ephrin-B1 as a repellent and attractant to control bidirectional branch extension in dorsal-ventral retinotopic mapping, *Development* 130, 2407-2418
- McLaughlin, T., O’Leary, D. D. M. (2005) Molecular gradients and development of retino-

topic maps, *Annu. Rev. Neurosci.* 28, 327-355

Mehta, M. R., Quirk, M. C., Wilson, M. A. (2000), Experience-Dependent Asymmetric Shape of Hippocampal Receptive Fields, *Neuron* 25, 707-715

Meinhardt, H., Gierer, A. (2000), Pattern formation by local self-activation and lateral inhibition, *Bio. Essays* 22, 753-760

Meister, M., Wong. R. O., Baylor, D. A., Shatz, C. J. (1991), Synchronous bursts of action potentials in ganglion cells of the developing mammalian retina, *Science* 252, 939-943

Miller, K. D. (1994), A model for the development of simple cell receptive fields and the ordered arrangement of orientation columns through activity-dependent competition between ON- and OFF-center inputs, *J. Neurosci.* 14, 409-441

Miller, K. D. (1996), Synaptic Economics: Minireview Competition and Cooperation in Synaptic Plasticity, *Neuron* 17, 371-374

Miller, K. D., MacKay, D. J. C. (1994), The role of constraints in Hebbian learning, *Neural Comput.* 6, 100-126

Moore IV, B. D., Alitto, H. J., Usrey, W. M. (2005), Orientation Tuning, But Not Direction Selectivity, Is Invariant to Temporal Frequency in Primary Visual Cortex, *J. Neurophysiol.* 94, 1336-1345

Movshon, J. A., Thomson, I. D., Tolhurst, D. J. (1978), Spatial summation in the receptive fields of simple cells in the cat's striate cortex, *J. Physiol.* 283, 53-77

Nelson, S. B., Turrigiano, G. G. (2008), Strength through Diversity, *Neuron* 60, 477-482

Nevian, T., Sakmann, B. (2006), Spine Ca^{2+} Signaling in Spike-Timing-Dependent Plasticity, *J. Neurosci.* 26, 11001-11013

Ngezahayo, A., Schachner, M., Artola, A. (2000), Synaptic Activity Modulates the Induction of Bidirectional Synaptic Changes in Adult Mouse Hippocampus, *J. Neurosci.* 20,

Nicol, X., Voyatzis, S., Muzerelle, A., Narboux-Nême, N., Südhof, T. C., Miles, R., Gaspar, P. (2007), cAMP oscillations and retinal activity are permissive for ephrin signalling during the establishment of the retinotopic map, *Nat. Neurosci.* 10, 340-347

Nishiyama, M., Hong, K., Mikoshiba, K., Poo, M. M., Kato, K. (2000), Calcium stores regulate the polarity and input specificity of synaptic modification, *Nature* 408, 584-588

O'Connor, D. H., Wittenberg, G. M., Wang, S. S. H. (2005), Dissection of Bidirectional Synaptic Plasticity Into Saturable Unidirectional Processes, *J. Neurophysiol.* 94, 1565-1573

Ohshiro, T., Weliky, M. (2006), Simple fall-off pattern of correlated neural activity in the developing lateral geniculate nucleus, *Nat. Neurosci.* 9, 1541-1548

Oja, E. (1982), A simplified Neuron Model as a Principal Component Analyzer, *J. Math. Biology* 15, 267-273

Olshausen, B. A., Field, D. J. (2004), Sparse coding of sensory inputs, *Current Opinion in Neurobiology* 14, 481-487

Paoletti, P., Bellone, C., Zhou, Q. (2013), NMDA receptor subunit diversity: impact on receptor properties, synaptic plasticity and disease, *Nat. Rev. Neurosci.* 14, 383-400

Pfeiffenberger, C., Cutforth, T., Woods, G., Yamada, J., Rentera, R. C., Copenhagen, D. R., Flanagan, J. G., Feldheim, D. A. (2005), Ephrin-As and neural activity are required for eye-specific patterning during retinogeniculate mapping, *Nat. Neurosci.* 8, 1022-1027

Pfeiffenberger, C., Yamada, J., Feldheim, D. A. (2006), Ephrin-As and Patterned Retinal Activity Act Together in the Development of Topographic Maps in the Primary Visual System, *J. Neurosci.* 26, 12873-12884

Pitkow, X., Meister, M. (2012), Decorrelation and efficient coding by retinal ganglion cells, *Nat. Neurosci.* 15, 628-635

Priebe, N. J., Ferster, D. (2005), Direction Selectivity of Excitation and Inhibition in

Simple Cells of the Cat Primary Visual Cortex, *Neuron* 45, 133-145

Rao, R. P. N., Sejnowski, T. (2003), Complex Cell-like Direction Selectivity through Spike-Timing Dependent Plasticity, *IETE J. Research* 49, 97-111

Rashid, T., Upton, A. L., Blentic, A., Ciossek, T., Knöll, B., Thompson, I. D., Drescher, U. (2005), Opposing Gradients of Ephrin-As and EphA7 in the Superior Colliculus Are Essential for Topographic Mapping in the Mammalian Visual System, *Neuron* 47, 57-69

Reid, R. C., Alonso, J. M. (1995), Specificity of monosynaptic connections from thalamus to visual cortex, *Nature* 378, 281-284

Reid, R. C., Soodak, R. E., Shapley, R. M. (1987), Linear Mechanisms of Directional Selectivity in Simple Cells of Cat Striate Cortex, *Proc. Natl. Acad. Sci. USA* 84, 8740-8744

Reid, R. C., Soodak, R. E., Shapley, R. M. (1991), Directional Selectivity and Spatiotemporal Structure of Receptive Fields of Simple Cells in Cat Striate Cortex, *J. Neurophysiol.* 66, 505-529

Reinagel, P., Zador, A. M. (1999), Natural scene statistics at the centre of gaze, *Network: Comput. Neural Syst.* 10, 341-350

Richards, B. A., Voss, O. P., Akerman, C. J. (2010), GABAergic circuits control stimulus-instructed receptive field development in the optic tectum, *Nat. Neurosci.* 13, 1098-1106

Rubin, C. M., van der List, D. A., Ballesteros, J. M., Goloshchapov, A. V., Chalupa, L. M., Chapman, B (2011), Mouse mutants the nicotinic acetylcholine receptor $\beta 2$ subunit display changes in cell adhesion and neurodegeneration response genes, *PLoS One* 6(4): e18626

Rust, N. C., Schwartz, O., Movshon, J. A., Simoncelli, E. P. (2005), Spatiotemporal Elements of Macaque V1 Receptive Fields, *Neuron* 46, 945-956

Sabatini, B. L., Maravall, M., Svoboda, K. (2001), Ca^{2+} signaling in dendritic spines,

Current Opinion in Neurobiology 11, 349-356

Sabatini, B. L., Oertner, T. G., Svoboda, K. (2002), The Life Cycle of Ca^{2+} Ions in Dendritic Spines, *Neuron* 33, 439-452

Saul, A. B., Carras, P. L., Humphrey, A. L. (2005), Temporal Properties of Inputs to Direction-Selective Neurons in Monkey V1, *J. Neurophysiol.* 94, 282-294

Saul, A. B., Feidler, J. C. (2002), Development of response timing and direction selectivity in cat visual thalamus and cortex, *J. Neurosci.* 22, 9245-2955

Saul, A. B., Humphrey, A. L. (1990), Spatial and temporal response properties of lagged and nonlagged cells in cat lateral geniculate nucleus, *J. Neurophysiol.* 64, 206-224

Senn, W., Buchs, N. J. (2003), Spike-Based Synaptic Plasticity and the Emergence of Direction Selective Simple Cells: Mathematical Analysis, *J. Comp. Neurosci.* 14, 119-138

Seol, G. H., Ziburkus, J., Huang, S., Song, L., Kim, I. T., Takamiya, K., Huganir, R. L., Lee, H. K., Kirkwood, A. (2007), Neuromodulators Control the Polarity of Spike-Timing-Dependent Synaptic Plasticity, *Neuron* 55, 919-929

Sernagor, E., Eglén, S. J., O'Donovan, M. J. (2000), Differential Effects of Acetylcholine and Glutamate Blockade on the Spatiotemporal Dynamics of Retinal Waves, *J. Neurosci.* 20, RC56 (1-6)

Sernagor, E., Young, C., Eglén, S. J. (2003), Developmental Modulation of Retinal Wave Dynamics: Shedding Light on the GABA Saga, *J. Neurosci.* 23, 7621-7629

Shah, R. D., Crair, M. C. (2008), Retinocollicular Synapse Maturation and Plasticity Are Regulated by Correlated Retinal Waves, *J. Neurosci.* 28, 292-303

Sherk, H., Stryker, M. P. (1976), Quantitative Study of Cortical Orientation Selectivity in Visually Inexperienced Kitten, *J. Neurophysiol.* 39, 63-70

Shi, J., Aamodt, S. M., Constantine-Paton, M. (1997), Temporal Correlations between Functional and Molecular Changes in NMDA Receptors and GABA Neurotransmission in

the Superior Colliculus, *J. Neurosci.* 17, 6264-6276

Shipton, O. A., Paulsen, O. (2014), GluN2A and GluN2B subunit-containing NMDA receptors in hippocampal plasticity, *Phil. Trans. R. Soc. B* 369, 20130163

Shon, A. P., Rao, R. P. N., Sejnowski, T. J. (2004), Motion detection and prediction through spike-timing dependent plasticity, *Network: Comput. Neural Syst.* 15, 179-198

Shouval, H. Z., Bear, M. F., Cooper, L. N. (2002), A unified model of NMDA receptor-dependent bidirectional synaptic plasticity, *Proc. Natl. Acad. Sci. USA* 99, 10831-10836

Simon, D. K., O'Leary, D. D. M. (1992a), Development of Topographic Order in the Mammalian Retinocollicular Projection, *J. Neurosci.* 12, 1212-1232

Simon, D. K., O'Leary, D. D. M. (1992b), Responses of Retinal Axons In Vivo and In Vitro to Position-Encoding Molecules in the Embryonic Superior Colliculus, *Neuron* 9, 977-989

Simpson, H. D., Mortimer, D., Goodhill, G. J. (2009), Theoretical Models of Neural Circuit Development, In: *Current Topics in Developmental Biology* (Hobert O., ed), Volume 87, pp. 1-51. Academic Press

Sjöström, P. J., Häusser, M. (2006), A Cooperative Switch Determines the Sign of Synaptic Plasticity in Distal Dendrites of Neocortical Pyramidal Neurons, *Neuron* 51, 227-238

Sjöström, P. J., Turrigiano, G. G., Nelson, S. B. (2001), Rate, Timing, and Cooperativity Jointly Determine Cortical Synaptic Plasticity, *Neuron* 32, 1149-1164

Sjöström, P. J., Turrigiano, G. G., Nelson, S. B. (2003), Neocortical LTD via Coincident Activation of Presynaptic NMDA and Cannabinoid Receptors, *Neuron* 39, 641-654

Smyth, D., Willmore, B., Baker, G. E., Thompson, I. D., Tolhurst, D. J. (2003), The Receptive-Field Organization of Simple Cells in Primary Visual Cortex of Ferrets under Natural Scene Stimulation, *J. Neurosci.* 23, 4746-4759

- Song, S., Abbott, L. F. (2001), Cortical Development and Remapping through Spike Timing-Dependent Plasticity, *Neuron* 32, 339-350
- Song, S., Miller, K. D., Abbott, L. F. (2000), Competitive Hebbian learning through spike-timing-dependent synaptic plasticity, *Nat. Neurosci.* 3, 919-926
- Sperry, R. W. (1963), Chemoaffinity in the orderly growth of nerve fiber patterns and connectins, *Proc. Natl. Acad. Sci. USA* 50, 703-710
- Stafford, B. K., Sher, A., Litke, A. M., Feldheim, D. A. (2009), Spatial-Temporal Patterns of Retinal Waves Underlying Activity-Dependent Refinement of Retinofugal Projections, *Neuron* 64, 200-212
- Sterratt, D. C., Lyngholm, D., Willshaw, D. J., Thompson, I. D. (2013), Standard Anatomical and Visual Space for the Mouse Retina: Computational Reconstruction and Transformation of Flattened Retinae with the Retistruct Package, *PLOS Comp. Biol.* 9(2): e1002921
- Stryker, M. P., Sherk, H., Leventhal, A. G., Hirsch, H. V. B. (1978), Physiological Consequences for the Cat's Visual Cortex of Effectively Restricting Early Visual Experience With Oriented Contours, *J. Neurophysiol.* 41, 896-909
- Sun, C., Warland, D. K., Ballesteros, J. M., van der List, D., Chalupa, L. M. (2008a), Retinal waves in mice lacking the $\beta 2$ subunit of the nicotinic acetylcholine receptor, *Proc. Natl. Acad. Sci. USA* 105, 13638-13643
- Sun, C., Speer, C. M., Wang, G. Y., Chapman, B., Chalupa, L. M. (2008b), Epibatidine Application In Vitro Blocks Retinal Waves Without Silencing All Retinal Ganglion Cell Action Potentials in Developing Retina of the Mouse and Ferret, *J. Neurophysiol.* 100, 3253-3263
- Syed, M. M., Lee, S., Zheng, J., Zhou, Z. J. (2004), Stage-dependent dynamics and modulation of spontaneous waves in the developing rabbit retina, *J. Physiol.* 560, 533-549
- Tavazoie, S. H., Reid, R. C. (2000), Diverse receptive fields in the lateral geniculate nucleus during thalamocortical development, *Nat. Neurosci.* 3, 608-616

- Thorn, F., Gollender, M., Erickson, P. (1975), The development of the kitten's visual optics, *Vision Res.* 16, 1145-1149
- Torborg, C. L., Feller, M. B. (2005), Spontaneous patterned retinal activity and the refinement of retinal projections, *Prog. Neurobiol.* 76, 213-235
- Turing, A. M. (1952), The Chemical Basis of Morphogenesis, *Phil. Trans. R. Soc. B* 237, 37-72
- Turrigiano, G. G., Nelson, S. B. (2004), Homeostatic plasticity in the developing nervous system, *Nat. Rev. Neurosci.* 5, 97-107
- Usrey, W. M., Alonso, J. M., Reid, R. C. (2000), Synaptic Interactions between Thalamic Inputs to Simple Cells in Cat Visual Cortex, *J. Neurosci.* 20, 5461-5467
- Wang, L., Rangarajan, K. V., Lawhn-Heath, C. A., Sarnaik, R., Wang, B. S., Liu, X., Cang, J. (2009), Direction-Specific Disruption of Subcortical Visual Behaviour and Receptive Fields in Mice Lacking the $\beta 2$ Subunit of Nicotinic Acetylcholine Receptor, *J. Neurosci.* 29, 12909-12918
- Warland, D. K., Huberman, A. D., Chalupa, L. M. (2006), Dynamics of Spontaneous Activity in the Fetal Macaque Retina during Development of Retinogeniculate Pathways, *J. Neurosci.* 26, 5190-5197
- Watt, A. J., Cuntz, H., Mori, M., Nusser, Z., Sjöström, P. J., Häusser, M., (2009), Traveling waves in developing cerebellar cortex mediated by asymmetrical Purkinje cell connectivity, *Nat. Neurosci.* 12, 463-473
- Weliky, M., Katz, L. C. (1997), Disruption of orientation tuning in visual cortex by artificially correlated neuronal activity, *Nature* 386, 680-685
- Wenisch, O. G., Noll, J., van Hemmen, J. L. (2005), Spontaneously emerging direction selectivity maps in visual cortex through STDP, *Biol. Cybern.* 93, 239-247
- White, L. E., Coppola, D. M., Fitzpatrick, D. (2001), The contribution of sensory experience to the maturation of orientation selectivity in ferret visual cortex, *Nature* 411, 1049-1052

- Wiesel, T. N., Hubel, D. H. (1963), Effects of visual deprivation on morphology and physiology of cells in the cat's lateral geniculate body, *J. Neurophys.* 26, 978-993
- Willshaw, D. J., von der Malsburg, C. (1976), How Patterned Neural Connections Can Be Set Up by Self-Organization, *Proc. R. Soc. Lond., B, Biol. Sci.* 194, 431-445
- Wimbauer, S., Wenisch, O. G., Miller, K. D., van Hemmen, J. L. (1997b), Development of spatiotemporal receptive fields of simple cells: I. Model formation, *Biol. Cybern.* 77, 453-461
- Wimbauer, S., Wenisch, O. G., van Hemmen, J. L., Miller, K. D. (1997a), Development of spatiotemporal receptive fields of simple cells: II. Simulation and analysis, *Biol. Cybern.* 77, 463-477
- Wittenberg, G. M., Wang, S. S. H. (2006), Malleability of Spike-Timing-Dependent Plasticity at the CA3-CA1 Synapse, *J. Neurosci.* 26, 6610-6617
- Wong, R. O. L., Meister, M., Shatz, C. J. (1993), Transient Period of Correlated Bursting Activity during Development of the Mammalian Retina, *Neuron* 11, 923-938
- Wong, R. O. L., Oakley, D. M. (1996), Changing patterns of spontaneous bursting activity in on and off retinal ganglion cells during development, *Neuron* 16, 1067-1095
- Xu, H. P., Furman, M., Mineur, Y. S., Chen, H., King, S. L., Zenisek, D., Zhou, Z. J., Butts, D. A., Tian, N., Picciotto, M. R., Crair, M. C. (2011), An instructive role for patterned spontaneous retinal activity in mouse visual map development, *Neuron* 70, 1115-1127
- Yates, P. A., Holub, A. D., McLaughlin, T., Sejnowski, T. J., O'Leary, D. D. M. (2004), Computational Modeling of Retinotopic Map Development to Define Contributions of EphA-EphrinA Gradients, Axon-Axon Interactions, and Patterned Activity, *J. Neurobiol.* 59, 95-113
- Yates, P. A., Roskies, A. L., McLaughlin, T., O'Leary, D. D. M. (2001), Topographic-Specific Axon Branching Controlled by Ephrin-As Is the Critical Event in Retinotectal Map Development, *J. Neurosci.* 21, 8548-8563

- Yu, X., Knierim, J. J., Lee, I., Shouval, H. Z. (2006), Simulating place field dynamics using spike timing-dependent plasticity, *Neurocomputing* 69, 1253-1259
- Zhang, J., Ackman, J. B., Xu, H. P., Crair, M. C. (2012), Visual map development depends on the temporal pattern of binocular activity in mice, *Nat. Neurosci.* 15, 298-307
- Zhang, L. I., Tao, H. W., Holt, C. E., Harris, W. A., Poo, M. M. (1998), A critical window for cooperation and competition among developing retinotectal synapses, *Nature* 395, 37-44
- Zheng, J., Lee, S., Zhou, Z. J. (2006), An transient network of intrinsically bursting starburst cells underlies the generation of retinal waves, *Nat. Neurosci.* 9, 363-371
- de Valois, R. L., Cottaris, N. P. (1998), Inputs to directionally selective simple cells in macaque striate cortex, *Proc. Natl. Acad. Sci. USA* 95, 14488-14493
- van Hooser, S. D., Li, Y., Christensson, M., Smith, G. B., White, L. E., Fitzpatrick, D. (2012), Initial neighbourhood biases and the quality of motion stimulation jointly influence the rapid emergence of direction preference in visual cortex, *J. Neurosci.* 32, 7258-7266
- von der Malsburg, C. (1973), Self-Organization of Orientation Sensitive Cells in the Striate Cortex, *Kybernetik* 14, 85-100