A novel simulation framework for modelling extracellular recordings in cortical tissue: implementation, validation, and application to gamma oscillations in mammals

Richard John Tomsett

Newcastle University

Faculty of Medical Sciences

Institute of Ageing and Health

Thesis submitted in partial fulfilment of the requirements for the degree of Doctor of Philosophy

March 2014

Abstract

This thesis concerns the simulation of local field potentials (LFPs) from cortical network activity; network gamma oscillations in particular. Alterations in gamma oscillation measurements are observed in many brain disorders. Understanding these measurements in terms of the underlying neuronal activity is crucial for developing effective therapies. Modelling can help to unravel the details of this relationship.

We first investigated a reduced compartmental neuron model for use in network simulations. We showed that reduced models containing <10 compartments could reproduce the LFP characteristics of the equivalent full-scale compartmental models to a reasonable degree of accuracy. Next, we created the Virtual Electrode Recording Tool for EXtracellular Potentials (VERTEX): a Matlab tool for simulating LFPs in large, spatially organised neuronal networks.

We used VERTEX to implement a large-scale neocortical slice model exhibiting gamma frequency oscillations under bath kainate application, an experimental preparation frequently used to investigate properties of gamma oscillations. We built the model based on currently available data on neocortical anatomy. By positioning a virtual electrode grid to match Utah array placement in experiments *in vitro*, we could make a meaningful direct comparison between simulated and experimentally recorded LFPs.

We next investigated the spatial properties of the LFP in more detail, using a smaller model of neocortical layer 2/3. We made several observations about the spatial features of the LFP that shed light on past experimental recordings: how gamma power and coherence decays away from an oscillating region, how layer thickness affects the LFP, which

neurons contribute most to the LFP signal, and how the LFP power scales with frequency at different model locations.

Finally, we discuss the relevance of our simulation results to experimental neuroscience. Our observations on the dominance of parvalbumin-expressing basket interneuron synapses on the LFP are of particular relevance to epilepsy and schizophrenia: changes in parvalbumin expression have been observed in both disorders. We suggest how our results could inform future experiments and aid in the interpretation of their results.

This thesis is dedicated to the Gods of Rock.

May they smile upon us always.

Acknowledgements

As a theoretical neuroscientist I am afraid to get my hands dirty in the lab, yet I am dependent on the discoveries of the experimentalists who are compelled to. Thanks are due, then, to our experimental collaborators, particularly Matt Ainsworth for kindly sharing his data. Many thanks also to Mark Cunningham and Miles Whittington, my co-supervisors, whose encyclopedic knowledge and willingness to help have been extremely valuable throughout my work.

The members of the Kaiserlab are owed a debt of gratitude, especially the unlucky few that shared an office with me and put up with both my creatively arranged desk space and the heavy metal leaking from my headphones. Special thanks must go to Jennifer Simonotto, who taught me the basics of digital signal processing, and who was always willing to help with my clueless questions. Extra special thanks go to Marcus Kaiser, my supervisor, for suggesting such interesting research directions, engaging in insightful discussions, and offering his unwavering support throughout.

Thank you to my family and friends, particularly those that have had the dubious pleasure of sharing the PhD experience with me in Newcastle upon Tyne, City of Dreams.

Finally, the greatest thanks are reserved for my parents, who have been unfailingly supportive and encouraging throughout my education, and Emma Kirkpatrick, who has not only been a saint to me during my write-up, but has also made the last four years extremely special.

Contents

List of Figures

List of Tables

Chapter 1

Introduction

This thesis is concerned with interpreting measurements of brain dynamics. Brains are fantastically complex organs, containing billions of electrically excitable neuronal cells communicating in intricate networks to generate global dynamics at a wide variety of spatial and temporal scales (Mitra and Bokil, 2008). Neuroscientists are faced with the task of explaining brain functions in terms of neuronal activity. However, measurements of this activity are beset with uncertainty: the more precise the measurement of the individual neurons, the fewer that can be measured simultaneously. This is problematic given that functionality results from the interaction of populations of neurons. Neuronal activity can be measured with various different techniques, but the relationship between these measurements and the behaviour of individual neurons is nontrivial (Brette and Destexhe, 2012).

Understanding brain activity measurements can be improved by developing models that explain the relationship between the measurements and the underlying neural activity. In this thesis, we develop tools and models for simulating extracellular electrode measurements in cortical networks. We then apply these tools to investigate the relationship between such measurements and neuronal network activity during network gamma oscillatory activity. We begin this chapter by expanding on the concept of brain dynamics and summarise how they can be measured. We then look at results linking brain dynamics, in particular gamma oscillations, to various cognitive and disease processes. Finally, we introduce the experimental and theoretical models that have been used to explain gamma oscillations, and describe how our work expands on these previous studies.

1.1 Brain dynamics

1.1.1 What are brain dynamics?

The study of a system's dynamics is the study of how the state of the system evolves in time. Brain dynamics, then, is the study of the evolution of the brain's state from one moment to the next; but "the brain's state" is an ill-defined concept that could be described by a near infinite combination of different variables. Investigations of brain dynamics are also limited by what can be measured in the living brain. Electrical activity is of particular interest, as measurements suggest that the brain uses electrical signals to encode stimuli and control muscle movement (Kandel et al., 2012) - in other words, an animal's behaviour is a direct result of the brain's electrical activity. Our research in particular is motivated by the study of electrical brain dynamics on timescales on the order of milliseconds to a few seconds. Electrical dynamics on this time scale are mostly driven by networks of neurons in the brain, so we also use the term neural dynamics to refer to the electrical dynamics of such networks, or neuronal dynamics to the activity of neurons.

1.1.2 Neural oscillations

Ever since the first electrical signals were measured from brains in the late 19th and early 20th centuries (Haas, 2003; Niedermeyer and Silva, 2005), neuroscientists have been fascinated by the oscillations they observed in their recordings. Over the past century, many different oscillatory regimes have been documented and correlated with cognitive processes or actions. Divided into frequency ranges, the more commonly studied of these are classified as delta (1.5 - 3Hz), theta (6 - 8Hz), alpha ($\sim$10Hz), beta (13 - 29Hz), gamma (30 - 80Hz) and very fast oscillations ($>$80Hz) (Buzsáki and Draguhn, 2004). Whether these oscillations are an epiphenomenon of neural dynamics or a fundamental mechanism of the brain's processing is still debated, but a wealth of evidence links characteristic oscillations with particular cognitive states or actions(Engel et al., 2001; Wang, 2010). Alterations in oscillatory activity have also been implicated in many brain diseases (Traub and Whittington, 2010). Clearly, even if oscillations are epiphenomena, the study of the mechanisms that cause them - and that cause pathological changes in them - will prove crucial

in efforts to prevent or treat these diseases (Traub and Whittington, 2010).

1.2 Measuring brain dynamics

Brain dynamics can be measured over many spatio-temporal scales by observing different variables using a range of modalities (Brette and Destexhe, 2012). Each method measures a different physical quantity that is in some way related to neuronal activity. Establishing this relationship for each modality is crucial for obtaining a better understanding of brain function and dysfunction, and provides us with our motivation for modelling extracellular recordings. Here we give a brief overview of the most common techniques for measuring network activity, placing our interest in the extracellular recordings in the broader context of general neural measurement techniques.

1.2.1 Global scale: fMRI, PET, EEG, MEG

Measurements of dynamics across the whole brain can be made using functional magnetic resonance imaging (fMRI), positron emission tomography (PET), electroencephalography (EEG) or magnetoencephalography (MEG) (Brette and Destexhe, 2012). fMRI and PET do not measure neural dynamics directly, but provide proxy measures based on variations in a quantity related to neural activity. In the case of fMRI this quantity is the oxygenation level of blood in the brain (Logothetis et al., 2001), and in the case of PET it is the level of a radioactive isotope indicating the concentration of glucose or oxygen (Raichle and Mintun, 2006) (though isotopes that bind to specific neuroreceptors can also be used, Kuhl et al., 1999; Catafau et al., 2010) . In each case, presence of glucose or of oxygenated blood is used as an indicator of the level of neural activity: higher neural activity uses more energy, so flow of oxygenated blood increases to areas of increased activity. While the spatial resolution of fMRI is extremely high - less than 1 mm^3 with modern scanners (Logothetis, 2008) - its temporal resolution is limited by the speed of the haemodynamic response, i.e. how rapidly oxygenated blood is routed to an area of high activity after that area becomes active. This response time is approximately two seconds (Logothetis and Wandell, 2004), and while temporal resolution can be improved by various methodological tricks, this physiological aspect is ultimately

the limiting factor in fMRI temporal resolution. The blood oxygenation level dependent (BOLD) signal is also difficult to relate to neural activity, as the details and function of the increased neural activity that gives rise to higher oxygenated blood flow are unknown. However, several studies have investigated how the BOLD signal is correlated with changes in extracellularly measured electric potentials (see below) . Interestingly, several results have suggested that the BOLD response correlates most strongly with the gamma frequency band (30 Hz - 80 Hz) of the local field potential (Koch et al., 2009; Niessing et al., 2005). Similar difficulties apply in PET imaging - additionally, PET can only be used to monitor brief tasks as the half-life of the radioactive isotopes is short. PET is frequently used in medical diagnosis as certain diseases show characteristic reductions in brain glucose usage (Jagust et al., 1991; Leiderman et al., 1992).

EEG and MEG both measure signals created directly by electrical neural activity: the electrical potential at the scalp in the case of EEG, and the magnetic field in the case of MEG (Nunez and Srinivasan, 2006). Both of these signals arise from the flow of currents in brain tissue, though each has different properties. The MEG signal is dominated by contributions from pyramidal neurons in the cortex's sulci, while the EEG contains contributions from the whole brain (though pyramidal neurons in the gyrii contribute most because of their relative proximity to the EEG electrodes, Lopes da Silva, 2013; Nunez and Srinivasan, 2006). This discrepancy is due to the nature of magnetic fields created by current flow. At a distance, current flow in pyramidal neurons can be approximated by current dipoles: current enters the neuron at one location and leaves at another. A current dipole produces a magnetic field according to Ampère's law. The direction of the magnetic field lines is given by the right-hand grip rule: they circle the axis of the current dipole. Pyramidal neurons positioned in gyrii have their current dipoles aligned approximately normally to the scalp, meaning that their magnetic fields are not visible outside the skull. In the sulcii, however, pyramidal current dipoles are aligned tangentially to the head surface, meaning that the magnetic field from these dipoles can be detected at the scalp (Ahlfors et al., 2010; Lopes da Silva, 2013). The greater specificity of the source locations picked up by MEG result in improved spatial resolution due to lower spatial smearing of the signal (Hämäläinen et al., 1993; Lopes da Silva, 2013). Skull and scalp also influence MEG less than EEG, further improving its spatial

resolution. Temporal resolution of each technique is very high, though, so rapid neural dynamics can be observed.

The spatial resolution of EEG can be greatly improved by placing the electrodes under the skull and dura, directly onto the surface of the brain. This method is intracranial EEG (iEEG) or the electrocorticogram (ECoG) (Strong et al., 2002). This method eliminates any filtering from the skull and scalp and places the electrodes closer to the current sources, providing much improved spatial resolution. ECoG is highly invasive, so can only be used in patients with intractable epilepsy to help surgeons localise the epileptic focal point they want to resect .

While both EEG and MEG are a direct result of neural activity, it is difficult to infer precisely what this neural activity is, and where it is located, given the recordings. Making such an inference requires solving an inverse problem: given a finite set of measurements, where are the dipole current sources located, and how large are they? Unfortunately, inverse problems are extremely difficult to solve as an infinite configuration of sources can result in the same set of measurements (Darvas et al., 2004). Educated guesses must be made, based on prior knowledge of brain shape and neural activity. Forward models have proved useful in testing inverse solutions: instead of inferring a model from data, we start with a known model and infer measurements based on our knowledge of the physics of the signal propagation (Mosher et al., 1999).

1.2.2 Local scale: VSD and ISOI

Studies on the local scale - in the region of millimetres or centimetres - can provide more detailed information about the local network activity, though such methods are currently limited to animal studies. Several techniques are available that operate on this scale, each working on different principles, thus providing different information about the operation of the network. We first briefly survey two optical imaging techniques: voltage-sensitive dyes (VSD), and intrinsic signal optical imaging (ISOI). These methods are primarily used for recording population activity at high spatial resolution.

In VSD imaging, a dye is applied to the neural tissue that binds to the surfaces of cell membranes. When the dye molecules are excited using monochromatic light, they fluoresce in proportion to membrane potential changes. This fluorescence is

detected by a highly sensitive camera. This signal is very small (approximately 1 percent fluorescence change Chemla and Chavane, 2012), so experiments must be repeated many times and averaged to obtain robust results. VSD imaging can be used both *in vitro* and *in vivo*, including in awake-behaving animals (Ferezou et al., 2006) (Grinvald and Hildesheim, 2004). It has very high temporal resolution, as the fluorescent properties of the dye change instantly with the cells' membrane potentials. Scattering of the emitted signal limits spatial resolution (Chemla and Chavane, 2010b), as does the use of a 2D sensor to image the sum of activity from the 3D cortical structure (Chemla and Chavane, 2010a). Individual neuronal contributions cannot be resolved; the signal at one camera pixel represents a weighted sum of the membrane potentials of all cell membranes under that pixel.

ISOI measures the amount of light reflected from brain tissue with no extra dyes added. Changes in the reflective properties of brain tissue that correlate with neural activity cause this reflected light signal to vary in time, giving a measure of underlying activity (Grinvald et al., 1986). In fact, some of the signal measured by VSD imaging comes from the intrinsic optical signal. The intrinsic optical signal has a relatively slow temporal resolution, as the changes in intrinsic tissue reflectiveness are slow. However, using this ISOI is less invasive than VSD or calcium imaging: no dye is required, so dye toxicity is not a problem. ISOI can also be used chronically in rat *in vivo*, as only a thinning of the skull is required to obtain a usable signal (Frostig and Chen-Bee, 2012).

1.2.3 Local scale: extracellular electrode recordings

Our focus in this thesis is on measurements made with extracellular electrodes (electrodes positioned in the space between neurons). These measurements can be used to study local network processing both *in vivo* and *in vitro*, in many different species. Electrodes positioned in this way record the extracellular potential at the electrode tip. This potential is created in exactly the same way as the EEG signal: neuronal membrane currents flowing in the resistive extracellular medium give rise to an electric potential (Buzsáki et al., 2012). However, the tip of an extracellular electrode is extremely small compared to an EEG electrode, so the number of neurons contributing to the extracellular potential is much smaller than the number contributing to the EEG (the contribution of a current source to the extracellular potential is

inversely proportional to distance, Nunez and Srinivasan, 2006). Extracellular electrodes therefore provide a measure of the local population activity on a similar scale to the imaging methods mentioned in the previous section.

Fluctuations in the extracellular potential are caused by the dynamics of neuronal membrane currents, resulting in measurements with a broad frequency range. This broadband signal is often split into high and low frequency bands. The high frequency (usually >500Hz, Pettersen et al., 2012) part of the signal is considered primarily to reflect the spike output of individual neurons very close to the electrode (Adrian and Moruzzi, 1939). This is termed the multi-unit activity (MUA), and can be used to obtain single neuron spike trains from extracellular recordings through spike-sorting methods (Pedreira et al., 2012; Quian Quiroga et al., 2004; Quian Quiroga, 2007). Individual spikes are resolvable because they are very rapid events ($\sim$1 ms) so often do not overlap in time, even at high firing rates. The low frequencies (usually <300Hz, Logothetis, 2003), called the local field potential (LFP), contain information about slower currents; it is this part of the signal that we are particularly concerned with in this thesis. A large proportion of the LFP is due to synaptic activity. Synaptic time constants can range from $\sim$1 ms to hundreds of milliseconds, so synaptic contributions to the LFP are likely to overlap in time, especially as there are thousands of times more synapses in the cortex than neurons (Shepherd, 2003). Additionally, the LFP contains contributions from slow membrane currents, for example from the hyperpolarising current in a neuron after it spikes or from sub-threshold membrane oscillations (Logothetis and Wandell, 2004). Consequently, interpreting the LFP in terms of neuronal activity is more difficult than interpreting the MUA, which provides some of the motivation for our modelling studies. The simplified assumption that the MUA contains information about spiking and the LFP contains information about synaptic activity has led to the characterisation of the MUA representing network output and the LFP representing network input (Logothetis, 2003), though as the majority of synapses in the cortex are local (Binzegger et al., 2004; Thomson et al., 2002), "output" should also be understood to mean internal recurrent processing. The ability to simultaneously record unit activity and LFPs at the same location is an important strength of extracellular electrode recordings: individual neuron firing can be observed in relation to a measure of the overall local network activity.

Recent increased interest in extracellular electrode recordings has been stimulated by the development of multi-electrode arrays (MEAs) (Buzsáki, 2004; Einevoll et al., 2012, 2013; Stevenson and Kording, 2011). MEAs are currently the only invasive recording method approved for use in humans (House et al., 2006), where applications include monitoring neural activity in patients with intractable epilepsy (Schevon et al., 2009, 2008) and recording signals to control brain-machine interfaces (BMIs) (Ball et al., 2004). Both of these applications would benefit from an improved understanding of LFP generation. In epilepsy, better knowledge of the LFP's spatial properties could help in more precise identification of seizure foci. In BMIs, further knowledge of the neuronal activity creating particular LFP features could improve learning of the control signal.

1.2.4 Single cell recordings

Single cell membrane potentials and currents can be measured using glass microelectrodes or patch electrodes. These recording methods are rarely used to measure network dynamics, but they can be used in conjunction with extracellular electrode recordings to provide more information on the sources of membrane currents in single neurons (as in Ainsworth et al., 2011, for example). Potential differences between the intracellular and extracellular media are measured by inserting a glass microelectrode filled with conductive medium into a neuron, taking care to cause minimal membrane damage, and a reference electrode into the extracellular space. The membrane potential can then be measured under different conditions. Intracellular electrodes can also be used to inject current into neurons to investigate their passive membrane properties. Patch-electrodes, which have a larger diameter tip than glass microelectrodes, do not pierce the cell membrane but attach to it by suction. This creates an extremely high resistance seal between the electrode and the membrane, allowing the recording of membrane currents with very high signal to noise ratio (Weckstrom, 2010). Currents through single ion channels can be measured, or, if the membrane is ruptured through further suction or using perforated patch techniques, whole-cell recordings of the membrane current can be performed. Single cell recordings using precursors of these techniques led to the development of Hodgkin and Huxley's famous model of action potential generation in the squid giant axon (Hodgkin and Huxley, 1952), which still underpins our understanding of

single neuron excitability today (Johnston and Wu, 1995; Koch, 1999).

Studies using these techniques have revealed the properties of the tens of different kinds of ion channel present in neuronal membranes (Hille, 2001), as well as membrane current and potential responses to different input patterns. Local functional networks have been mapped out by injecting current into a presynaptic neuron, and recording membrane responses from a series of potential postsynaptic neurons, so that connectivity statistics can be quantified (Thomson and Lamy, 2007; Thomson et al., 2002; Lefort et al., 2009).

One limitation of intracellular recordings is the extreme difficulty of recording from thinner neurites, meaning that the membrane potential or current can usually only be recorded at the soma (though recording at the apical dendrite of pyramidal neurons has been done (Hay et al., 2011), these studies are the exception rather than the norm). This is not a problem for many experimental questions, but means that we suffer from a lack of data about current flow in dendrites. Very high-density MEA recordings may improve this situation: if an MEA such as the 10 000 electrode array described in (Frey et al., 2009) is positioned over a neuron, and that neuron is stimulated separately, then the extracellular potential over the extent of that neuron's dendritic tree can be measured. As we know the theoretical relationship between neuronal membrane currents and extracellular potentials (see Chapter 2), it may be possible to use this data to further constrain the active membrane conductance parameters in that neuron's dendrites (Gold et al., 2007).

1.3 Relating animal behaviour to measured dynamics

The reason for measuring brain dynamics is to elucidate the mechanisms by which animal behaviours are implemented. It is hoped that knowledge of these mechanisms will improve the human condition in three ways: medically, technologically, and philosophically. On the final point, it is believed that a knowledge of the brain will help us to understand ourselves better - why we are as we are. On the second point, it is hoped that knowledge of the how particular behaviours are implemented by nature will allow us to construct better machines that exhibit these behaviours (for example, object recognition or motor control). On the first point, it is suspected

that knowing more about mechanisms in the brain will allow us to create better medicines and therapies for when our brains break.

The study of gamma oscillations pertains to all three of these points. Gamma rhythms have been observed in many brain areas as a correlate of an assortment of high-level cognitive processes (Buzsáki, 2006). They have been proposed as a mechanism by which the brain encodes and transmits information between brain areas (Ainsworth et al., 2012), as well as facilitating learning through synaptic plasticity (Osipova et al., 2006). Finally, abnormalities in gamma oscillatory activity have been observed in a variety of brain disorders (Başar, 2013; Cunningham et al., 2006; Koenig et al., 2005; Lee et al., 2003a). Understanding precisely how gamma oscillations are generated by healthy brains may provide important clues about disease mechanisms and, potentially, how to counteract them.

1.3.1 Gamma rhythms and cognition

The prevalence of gamma frequency activity across many brain regions, and experimental correlations with various tasks and cognitive states, has led to a wide variety of different functions being ascribed to gamma rhythms. The focus of most studies in humans has been on neocortical gamma rhythms, as these are easiest to record with EEG. Gamma-band power increases have been measured by EEG, ECoG and MEG in humans during working memory tasks (Mainy et al., 2007; Tallon-Baudry et al., 2005, 1998), and during selective attention to stimuli (Bauer et al., 2006). Increases in gamma-band power and synchronisation between brain areas were observed (using ECoG) in visuomotor task performance in human sensorimotor cortex (Aoki et al., 1999). Long term memory formation is predicted by the gamma activity measured during the encoding phase (Gruber et al., 2004; Jutras et al., 2009; Osipova et al., 2006; Sederberg et al., 2007).

The theory of why gamma oscillations are involved in this variety of processes relates to the properties of information processing during gamma activity. Gamma oscillations were proposed as the physical method through which the brain binds object representations together in the late 1980s/early 1990s (Eckhorn et al., 1990, 1988; Engel et al., 1991). Since then, a comprehensive theory of how gamma oscillations are suited to implementing binding has been developed. Pascal Fries (Fries, 2009) proposes that gamma oscillations are the mechanism by which groups of neur-

ons functionally segment converging inputs, and as such are a fundamental computational process in the cortex. Critics have argued that, as gamma is recorded during so many cognitive processes, it may not be serving any particular function, but rather represents an epiphenomenon of neural network processing (Burns et al., 2011; Herrmann et al., 2004; Ray and Maunsell, 2010). However, while the gamma frequency may not be directly used by the brain for information encoding or stimulus binding as hypothesised, its pervasiveness means that understanding the mechanisms that give rise to gamma rhythms will likely help us to understand the general principles of electrical dynamics in neuronal networks.

1.3.2 Gamma oscillations and disease

Just as gamma oscillations are correlated with cognitive states, alterations to gamma oscillations have been observed in patients with brain disorders (Herrmann and Demiralp, 2005). Decreases in gamma power or phase synchrony have been found in patients with Alzheimer's disease (Koenig et al., 2005; Stam et al., 2002), autism and Williams syndrome (Grice et al., 2001), as well as after stroke and brain injury (Molnár et al., 1997; Slewa-Younan et al., 2002) and with increasing age (Böttger et al., 2002).

The picture is somewhat more complicated in schizophrenia. Gamma-band power and synchronisation across areas is, in general, reduced for spontaneous, task-related, and sensory-evoked gamma oscillations (Gallinat et al., 2004; Green et al., 2003; Haig et al., 2000; Kissler et al., 2000; Light et al., 2006). However, the degree of synchrony in the gamma band correlates with the positive symptoms of schizophrenia (e.g. hallucinations) (Lee et al., 2003b; Spencer et al., 2004). Positive symptoms also correlate with increased gamma power during sleep in schizophrenic patients (Tekell et al., 2005). A recent review is provided in (Sun et al., 2011).

In epilepsy, increases in gamma-band activity related to seizures are often seen. Gamma power is frequently observed to increase before seizure onset (Alarcon et al., 1995; Fisher et al., 1992). Gamma activity can also show greatly increased amplitude during the interictal (between-seizure) phase in patients with primary generalised epilepsy (Willoughby et al., 2003). Increases in gamma activity have been observed close to the seizure focal point in neocortical epilepsy (Worrell et al., 2004).

The relationship between schizophrenia and epilepsy is not clear, though a re-

cent study uncovered a significantly increased incidence of schizophrenia in epilepsy patients (adjusted hazard ratio of 7.65) and of epilepsy in schizophrenia patients (adjusted hazard ratio of 5.88) in a retrospective cohort analysis (Chang et al., 2011). However, we find it relevant that a reduction in the calcium buffering protein parvalbumin (PV) is apparent in both diseases (Andrioli et al., 2007; Beasley and Reynolds, 1997; DeFelipe et al., 1993; Hashimoto et al., 2003; Marco and DeFelipe, 1997). PV is expressed in basket interneurons (DeFelipe, 1997), a fast spiking interneuron type that targets the perisomatic region of pyramidal neurons. These interneurons are crucial for the generation of gamma oscillations in neocortex and hippocampus (Bartos et al., 2007), so it is no surprise that reductions in PV affect gamma activity. Establishing why the effect on gamma oscillations is different in each condition could help to unravel the links between these disorders.

1.3.3 Gamma oscillation models

The interest in finding the mechanisms of gamma oscillations has led to the development of various *in vitro* experimental preparations that allow the gamma-generating network to be studied in isolation. Most of these preparations have looked at gamma activity in hippocampus (Atallah and Scanziani, 2009; Buhl et al., 2003; Draguhn et al., 1998; Ferguson et al., 2013; Oren et al., 2006; Traub et al., 1996a; van der Linden et al., 1999; Whittington et al., 2001, 1997b,a; Traub et al., 2003b; Whittington et al., 1995), entorhinal cortex (Cunningham et al., 2003, 2004a; Dickson et al., 2000; van der Linden et al., 1999) and neocortex (Ainsworth et al., 2011; Cunningham et al., 2004b; Haenschel et al., 2000; Traub et al., 2005a; Whittington et al., 1997b, 1995). Early studies showed how networks of tonically driven interneurons could generate a gamma-frequency rhythm, measured in the LFP in rat hippocampal slices, when synaptic excitation was blocked (Whittington et al., 1995). The inhibitory rhythm generated in this way - interneuron network gamma (ING) - was shown to be able to entrain the outputs of the pyramidal neurons to the oscillation. This model was later shown also to be dependent on the presence of gap junctions between inhibitory neuron dendrites, which cause the interneurons to track changes in each others' membrane potentials (Traub et al., 2001a). ING rhythms only last for a few seconds.

The tetanic model of gamma oscillations (where gamma rhythms are induced

by high frequency electrical stimulation) also lasts for a short time, but in this model excitatory synapses are not blocked; in fact the oscillation depends on them (Whittington et al., 1997a). Pyramidal neurons and interneurons fire on almost all oscillation cycles, so this mode of operation is referred to as pyramidal-interneuron network gamma (PING). PING rhythms capture many of the properties of sensory-evoked gamma oscillations (Traub et al., 1999a).

Another model referred to as "persistent gamma" - gamma rhythms that persist for hours rather than transiently after stimulation - was reported in (Fisahn et al., 1998) to occur when carbachol was applied in the bath to hippocampal area CA3. An important feature of this type of oscillation is that pyramidal neurons do not fire on every oscillation cycle, but more sparsely, as is often the case *in vivo*. The robustness of this oscillation allowed a huge variety of experiments to be performed to determine the relevant mechanisms (Traub and Whittington, 2010). Persistent gamma oscillations are dependent on the interneurons receiving phasic excitation (Fuchs et al., 2007), but are not dependent on pyramidal-pyramidal excitation (Fisahn et al., 1998). Persistent gamma also appears to depend on the presence of gap junctions, this time between the axons of pyramidal cells rather than between interneurons (Hormuzdi et al., 2001) (though gap junctions between interneurons modulate the gamma power (Buhl et al., 2003)). It is currently unknown exactly how the *in vitro* persistent gamma model relates to gamma oscillations with sparse pyramidal cell firing *in vivo*. We return to the mechanisms of persistent gamma generation in Chapter 5.

1.4 Theoretical oscillation models

Computational models have been crucial in developing the theory and guiding experimental studies into gamma oscillations. In this section, we review previous work on theoretical models of neural oscillations more generally, ranging from the highly abstract to the fantastically detailed, and draw some conclusions about how to construct suitable models to investigate our particular area of interest.

1.4.1 Simplified oscillation models

The mathematical study of oscillations in physical, chemical and biological systems has a long history, going back to studies of the simple harmonic motion of a single pendulum or spring. In neuroscience, electrical oscillations in brains were noted as early as the late nineteenth century, and the invention of the electroencephalogram (EEG) in the 1920s allowed oscillations to be observed in humans from scalp recordings (Niedermeyer and Silva, 2005). The underlying mechanisms of these oscillations are quite different from those of simple linear systems, involving many complex, non-linear interactions at multiple temporal and spatial scales, but surprisingly (or perhaps not so surprisingly (Goldenfeld and Kadanoff, 1999)) many of their features can be described mathematically by relatively few equations. Hordes of researchers have studied oscillatory phenomena in biological systems and a comprehensive overview is beyond the scope of this thesis, though it is worth mentioning a few notable applications of mathematics to study biological oscillators. Art Winfree was a pioneer in this field, as well as mathematical biology in general; his seminal 1967 paper in the Journal of Theoretical Biology (Winfree, 1967) helped lay the foundations for much future work on biological rhythms. He studied models of coupled oscillators - naturally oscillating units interacting according to some non-linear coupling mechanism affecting the oscillation phase of coupled units. Winfree's generalised models are difficult to analyse analytically, leading to the development of simpler formulations, including the widely studied Kuramoto model (Kuramoto, 1975).

Since these early contributions, coupled oscillator theory has been applied to many problems in biology, not least in modelling rhythms in the brain. Individual neurons can be thought of as oscillators as they will spike periodically when they receive a tonic excitatory input. Several important results about the mathematics governing different aspects of synchronous behaviour in neural systems include coupled oscillators in central pattern generators for rhythmic motion (Kopell and Ermentrout, 1988), properties of coupled oscillator systems leading to variation in oscillation amplitude (Aronson et al., 1990), the conditions necessary for stable phase-locking between neuronal oscillators and mechanisms of frequency control (Gerstner et al., 1996; Kopell and Ermentrout, 2002), the role of inhibitory coupling in rhythmogenesis (Wang and Rinzel, 1992), mechanisms of coherent oscillations in heterogeneous oscillator networks (White et al., 1998), in networks with sparse, ran-

dom connectivity (Börgers and Kopell, 2003), in networks with noisy drive (Börgers and Kopell, 2005), in networks with delays (Roxin et al., 2005) and many other regimes.

While these studies have provided important results about the dynamics of oscillations under various conditions, coupled oscillator theory is often not enough to describe real oscillations occurring in neural tissue (Wang, 2010). Neurons modelled as coupled oscillators will fire on every cycle of the oscillation, whereas many oscillations observed *in vitro* and *in vivo* show cells firing more sparsely than the network oscillation frequency (Bragin et al., 1995; Buhl et al., 1998; Colgin et al., 2009; Csicsvari et al., 2003; Fisahn et al., 1998; Whittington et al., 2000).

Xiao-Jing Wang presents a different framework for understanding oscillatory activity, based on work by Brunel, Geisler, Hakim, Hansel and himself (Brunel, 2000; Brunel and Hakim, 1999, 2008; Brunel and Hansel, 2006; Brunel and Wang, 2003; Geisler et al., 2005) . In this approach, neurons generate Poisson-like spike trains driven by high background synaptic noise. Inhibitory networks of such neurons, coupled by connections including transmission delay times, will exhibit high frequency, sparse synchronisation. Neurons that happen to fire together in the noisy regime will hyperpolarise connected neurons synchronously; these neurons will then be more likely to fire close together, as the synaptic noise, on average, will raise them to their firing thresholds from their hyperpolarised states at a similar rate. As the network is randomly connected, the synchronising pulses are started by the noise in random subsets of neurons, and the tonic firing and noise are much stronger than the oscillation, the network can exhibit oscillations with sparse firing and individual spike trains almost uncorrelated with the network frequency (Wang, 2010). The network frequency can be shown analytically and in simulations to be independent of neuronal firing rate, with synaptic and individual cell membrane properties shaping the frequency of the network rhythm (Brunel and Wang, 2003).

1.4.2 Detailed network models

As more and more experimental data on individual neuron structure and behaviour has been collected, increasingly complex models have been investigated that attempt to incorporate as much detail as is practical. This approach is in contrast to that described previously, where the simplest possible model to describe a phenomenon

is used. Each approach has its own benefits and drawbacks. Reduced models are more readily analysed mathematically, allowing proofs to be constructed about their behaviours and intuitive insights to be gained about the system's behaviours, while detailed models must be simulated, introducing problems in interpretation and analysis, as well as difficulties with numerical methods and computational complexity. The more detailed a model, the more parameters there are to adjust, leading to many possible parameter combinations that create similar behaviours in the model, but that may not be biologically correct (Prinz et al., 2004). However, by carefully incorporating biological detail from experimental data rather than creating abstract phenomenological descriptions, we can be more confident that a model's behaviour represents what is occurring biologically rather than a potentially unhelpful caricature of the system. Each variable and parameter also has a direct biological correlate, which makes model results more easily relatable to experimental data.

Perhaps the best example of the use of detailed models in investigating neuronal network oscillations, and of the interaction between modelling and experiment, is the work of Roger Traub and colleagues (for overviews, see (Traub et al., 2005b, 1999a; Traub and Whittington, 2010)). Over many years Traub et al. have created extremely detailed models of individual neuron types in various brain areas, including hippocampus, thalamus and neocortex. These models reproduce the highly complex membrane properties of the different neuron types, as well as approximating their dendritic trees and axons, to create neuron models that behave very much like their real counterparts under a wide variety of conditions. These have been used in network models to study the detailed mechanisms affecting oscillatory activity in different brain areas. Some key results include a model of a hippocampal slice displaying emergent rhythms with properties similar to those observed *in vitro* (Traub et al., 1989), the role of excitatory currents on inhibitory interneurons in defining the characteristics of oscillations and an explanation for gamma frequency "tails" after synchronised activity bursts (Traub et al., 1996a), the firing of spike doublets (two spikes in quick succession) by interneurons in coordinating gamma oscillations over long distances (Traub et al., 1996b), the different synchronisation properties of different frequency oscillations (Kopell et al., 2000), a novel mechanism for synchrony at high frequencies by way of gap junctions between axons of excitatory cells (Traub et al., 1999b) and a possible role for such gap junctions

in generating seizures (Roopun et al., 2010; Traub et al., 2001b), the role of gap junctions between interneuron dendrites in enhancing gamma frequency coherence (Traub et al., 2001a), mechanisms inducing rhythmic bursting behaviour in superficial layer pyramidal neurons (Traub et al., 2003a) and a potential role for these neurons in gamma generation in these layers (Cunningham et al., 2004b; Traub et al., 2005a), and mechanisms of rhythm generation through period concatenation between different oscillation frequencies (Kramer et al., 2008; Roopun et al., 2008). These and many other studies suggest a more complicated picture of how the brain generates oscillations than those suggested by the simplified models discussed previously. Both kinds of model are useful for different questions, and can provide different types of insight into a system's behaviour. Choosing the "right" level of abstraction for a model is a difficult challenge for the prospective modeller (Herz et al., 2006; Staley et al., 2011)..

1.5 Contributions of this thesis

1.5.1 Models and tools for simulating LFPs in networks

Recent results have indicated the importance of using multi-compartment neuron models when modelling the LFP (Lindén et al., 2010, 2011; Pettersen and Einevoll, 2008; Pettersen et al., 2012). However, these models are computationally expensive, making network simulations unfeasible without access to large computing resources (Markram, 2006). Most current models of neuronal network activity usually use simplified, single-compartment neurons and homogeneous networks, so various attempts have been made to estimate a realistic LFP signal from such models to (Einevoll et al., 2013). Recent studies by Mazzoni et al. (Mazzoni et al., 2008, 2010) have successfully captured features of experimentally recorded LFPs without considering dendritic morphology, tissue filtering or spatial location. They simulated a simplified model with two populations (excitatory pyramidal cells and inhibitory interneurons) of single compartment leaky integrate-and-fire neurons, calculating the LFP as the sum of the absolute values of excitatory and inhibitory currents on the pyramidal neurons only. This summation method was further developed in (beim Graben and Rodrigues, 2013) in an effort to incorporate knowledge about the filtering effects of

pyramidal neuron dendrites. However, none of these methods considers the spatial aspects of LFP generation, which are crucial when predicting the LFP at different points in spatially organised models.

We therefore investigated a reduced compartmental model, assessing its suitability for producing realistic spatial LFP characteristics and the improvement in simulation time. We also developed a new simulation tool for the efficient simulation of LFPs in spatially organised cortical networks using the reduced compartmental models. We implemented this tool in Matlab (Mathworks Inc., Natick, MA, USA) and provide a simple interface to maximise the tool's utility to the neuroscience community.

1.5.2 Model of LFP generation *in vitro* during persistent gamma

High density multi-electrode arrays are increasingly being used to record from several local network sites simultaneously. This necessitates the development of spatially organised network models that incorporate experimental anatomical knowledge in order to explain variations in activity observed in the experimental recordings. Our experimental collaborators at the Newcastle University Institute of Neuroscience have been using the Utah multi-electrode array to record from cortical slices *in vitro* (Ainsworth, 2013). This MEA is a 10 by 10 grid with a 0.4 mm spacing between electrodes; it covers a 3.6 mm by 3.6 mm square. The Utah array can therefore sample extracellular potentials from the full depth of the cortex, as well as over a substantial lateral distance. We used our simulation tool to implement a model of neuronal dynamics in a neocortical slice preparation to use in conjunction with the Utah array experiments. The model is large enough to include a virtual Utah array, and incorporates details from the literature regarding neuronal composition, spatial organisation, connectivity and dynamics. We simulate gamma oscillations in our model, and compare simulated LFPs with experimental recordings.

1.5.3 Investigation of LFP features during gamma oscillations

We provide a thorough analysis of the spatial properties of localised gamma oscillations, inspired by the observation of "patches" of gamma power in the LFP *in vitro* (Ainsworth, 2013). Considering only neocortical layer 2/3 - the gamma-oscillation generating layer - to simplify the analysis, we generated localised gamma oscillations to study the spatial distribution and decay of gamma power and the phase-inversion observed experimentally between layer 3 and layer 1 (Chrobak and Buzsáki, 1998; Cunningham et al., 2003; Dickson et al., 2000). We explain some previously confusing results regarding this phase inversion, and make some experimentally testable predictions. We also investigate how each neuron type contributes to the LFP during both random and gamma activity, something that is currently impossible to do experimentally. We make a surprising observation regarding the influence of soma-targeting basket interneurons on the LFP. We relate these findings to current research in schizophrenia and epilepsy, and make suggestions for future experimental and modelling studies based on our findings.

1.6 Summary

We are interested in understanding brain dynamics generated by the coordinated activity of neurons. To study brain dynamics one must measure brain dynamics. It is therefore crucial to understand precisely what is being measured in terms of the underlying neural activity. Establishing this relationship for each measurement modality is nontrivial, but understanding can be improved by building models that explain the measurements.

Our particular topic of interest is the local field potential generated during gamma oscillations. Gamma activity has been linked to many cognitive processes, and alterations to gamma rhythms are implicated in many brain disorders. Understanding the mechanisms of functional and dysfunctional gamma oscillations will provide important insights into the mechanisms of these disorders and potentially suggest avenues for improved treatments. Most work in gamma oscillations has been, and continues to be, performed using EEG or extracellular electrodes, with

an increasing use of multi-electrode arrays to study activity across space. Understanding the neuronal contributions to, and spatial properties of, gamma oscillations will greatly aid the interpretation of MEA recordings. We contribute to this understanding by studying models of the LFP generated by neuronal networks exhibiting gamma activity.

1.7 Thesis structure

The remainder of the thesis is separated into a further five chapters:

- In Chapter 2, we summarise the physics of extracellular potential generation in neural tissue and describe our experiments on LFPs generated by compartmental neuron models containing fewer than ten compartments (Bush and Sejnowski, 1993). Our focus was on simulating LFPs from network activity, so we needed a neuron model that would generate realistic extracellular potentials (Łęski et al., 2013; Lindén et al., 2011; Pettersen et al., 2012), but that was not too computationally demanding to simulate. We describe our experiments to test the reduced model's suitability for this.

- Chapter 3 introduces the Matlab-based software we created to simulate LFPs in network models: the Virtual Electrode Recording Tool for EXtracellular Potentials (VERTEX). We developed VERTEX as there was no currently available simulation tool for modelling extracellular potentials in neuronal networks. This chapter describes the mathematical models implemented in VERTEX, how we implemented these models to try to maximise the simulation speed for large populations, and the parallel programming methods used to further improve simulation speeds.

- Chapter 4 describes a large-scale neocortical slice model that we implemented in VERTEX. The idea behind this model was to incorporate anatomical and electrophysiological data into a model that represented the neurons in an equivalent experimental slice on a roughly 1:1 basis. We describe the data the model is based on, and the relevant parameters. We finally compared LFPs generated at 100 virtual electrode locations with Utah array recordings from

macaque neocortex *in vitro*. We describe the promise of this kind of large-scale model as well as the many difficulties we faced in its implementation.

- In Chapter 5, we develop a smaller model of neocortical layer 2/3 in order to investigate the properties of the LFP during gamma oscillations more thoroughly. We describe our experimental setup, in which we generate a localised gamma oscillation surrounded by randomly firing neurons. This allowed us to investigate how visible the gamma oscillation is outside of the generating region, how each neuron group contributes to the LFP signal in both gamma and desynchronised activity, and how gamma coherence decays across space. We also make some theoretical predictions regarding the effect of layer thickness on the LFP, and discuss the frequency scaling observed in our models in the context of the current debate in the literature regarding the origin of low-pass filtering of the LFP (Bédard et al., 2006b; Bédard and Destexhe, 2009; Bédard et al., 2004, 2010; Łęski et al., 2013; Lindén et al., 2010; Logothetis et al., 2007).

- In Chapter 6, we summarise our results and discuss them in the context of interpreting experimental LFP signals, particularly those recorded in human patients. We suggest several avenues for further work, including modelling recent data in intra-laminar phase synchrony during gamma oscillations *in vitro* and basket interneuron changes in schizophrenia and epilepsy.

Chapter 2

Modelling Local Field Potentials

2.1 Introduction

In this chapter, we summarise experimentally observed features of extracellular potentials, the physics of how extracellular potentials are generated, and how LFPs arising from neuronal activity can be modelled using compartmental neuron models. Previous modelling studies have used compartmental models created from whole-cell reconstructions (Einevoll et al., 2007; Gold et al., 2006; Gratiy et al., 2011; Holt and Koch, 1999; Łęski et al., 2013; Lindén et al., 2010, 2011; Pettersen and Einevoll, 2008; Reimann et al., 2013), but such compartmental models are computationally expensive to simulate, making them impractical for large network simulations. We identified a potentially suitable model reduction (Bush and Sejnowski, 1993) and compared LFPs simulated using the detailed models against their reduced versions, using the same methodology as (Lindén et al., 2011), to investigate the suitability of the reduced models for simulating LFPs.

2.2 What do extracellular electrodes measure?

Extracellular electrodes measure the electric potential in the extracellular space; that is, the space around the neurons. Extracellular electrodes can be used *in vivo* and *in vitro*, in many different brain areas and species. In section 2.2.1, we describe some experimentally observed features of extracellular potentials, and models of the processes that give rise to these features. We then outline the physics of extracellular potential generation from neuronal membrane currents.

2.2.1 Features of extracellular potentials

The extracellular potential is a time-varying signal that represents some measure of the dynamics of a local population of neurons around the measurement point. A crucial question, then, is what exactly is meant by "local": how far away can neurons be from the measurement point and still contribute significantly to the extracellular potential? Several groups have attempted to address this question experimentally; however, results have been inconsistent, with estimates ranging from several millimetres (Kajikawa and Schroeder, 2011; Kreiman et al., 2006; Logothetis et al., 2001; Mitzdorf, 1985) down to a few hundred microns (Berens et al., 2008; Katzner et al., 2009; Liu and Newsome, 2006; Xing et al., 2009). Differences in methodology explain to some extent why such discrepancies in results exist, but not at the level of the neuronal activity that might have differed between experiments.

Recent studies have attempted to address this question through modelling (Łęski et al., 2013; Lindén et al., 2011). The advantage of this approach is that the positions of the neurons and the neuronal activity is specified and can be systematically varied, so changes in the extracellular potential can be related directly to the known changes in the underlying activity. Lindén et al. (2011) studied the spatial spread of the extracellular potential created by synaptic activity in populations of compartmental neurons with realistic morphologies (Mainen and Sejnowski, 1996). They found that the spatial spread of the LFP varied depending on the location of the synapses on the dendrites and the level of correlation between the synaptic inputs. Correlated synaptic input increases both the magnitude and spread of the LFP, particularly when the input was localised to a particular dendritic region (e.g. only to the basal dendrites): as each neuron's contribution to the LFP sums linearly (see below), there is greater opportunity for constructive interference when input correlations are present. The spatial range also varied at different recording depths above and below the level of the neurons' soma positions: it was smallest at the soma level, and increased above and below this level. The magnitude of the signal was greatest at the soma level, and decreased above and below this level. Łęski et al. (2013) extended this approach, investigating the spatial spread of different frequency components of the LFP signal. They found that the spatial reach of higher frequencies was slightly lower than that of lower frequencies when neurons received uncorrelated input, and much lower when input correlations were introduced. These

studies provide a possible explanation for the discrepancies in experimental results: firstly, the position of the electrodes relative to the LFP generating populations in each experiment might have been inconsistent, and secondly, the synaptic input correlations could have been quite different (Einevoll et al., 2013). These modelling results show that considering the spatial organisation of the neuronal populations, positioning of the electrode, and the synaptic activity are crucial factors to consider when interpreting or predicting LFP measurements.

The results pertaining to frequency scaling reported in (Łęski et al., 2013) build on previous studies inspired by the experimental observation that extracellular potentials appear to be frequency filtered: fast extracellular potentials such as spikes are only visible from neurons within a very small distance from the electrode (Buzsáki, 2004), compared with the lower frequencies that are not only greater in amplitude but also visible from a greater distance. Some debate has occured regarding the origin of these frequency characteristics (Destexhe and Bedard, 2013; Einevoll et al., 2013). One possibility is that brain tissue acts as a low-pass filter. Experimental measurements have again provided conflicting results - (Gabriel et al., 1996) suggesting that this is the case, while (Logothetis et al., 2007) indicated that at physiologically relevant frequencies, the frequency filtering of brain tissue is negligible. Modelling results have suggested that the mechanism behind any tissue filtering effects could be ionic diffusion (Bédard and Destexhe, 2009; Bédard et al., 2010) and inhomogeneities in the extracellular medium (Bédard et al., 2006b, 2004). However, these models did not consider the frequency filtering effects due to passive current flow through the spatial extent of neurons' dendrites, which has been shown in other modelling studies to produce a low-pass frequency filtering effect as well (Łęski et al., 2013; Lindén et al., 2011; Pettersen and Einevoll, 2008). These two effects may both contribute to the overall frequency profile of the extracellular potential.

2.2.2 Physics of extracellular potentials

To understand the origin of the signal measured by an electrode inserted into brain tissue, we must consider the physics of electric fields in biological tissue. Biological tissues are considered as volume conductors: rather than the discrete elements of an electric circuit, the tissue is a continuum containing currents and potentials that are

functions of position (Plonsey and Barr, 2000). The equation relating the current flowing in the volume conductor to the electric potential is:

$$\nabla\sigma \cdot \nabla\Phi = -I_v, \tag{2.1}$$

(Nicholson and Freeman, 1975), where Φ is the electric potential, I_v is the current source density (or CSD; i.e. the volume density of current entering the extracellullar space (Pettersen et al., 2012)), and σ is the conductivity of brain tissue. If σ is not space-dependent - a reasonable assumption in many cases (though see discussion below) - this becomes Poisson's equation (Nicholson and Freeman, 1975; Plonsey and Barr, 2000):

$$\nabla^2\Phi = -\frac{I_v}{\sigma}. \tag{2.2}$$

This equation allows us to relate the electric potential, which is the quantity measured by the extracellular electrode, to the currents that create it. These currents are primarily neuronal membrane currents (Johnston and Wu, 1995); though membrane currents in other cell types, such as glia, will affect the potential, these currents are thought to be much smaller than those from neurons. If we know the tissue conductivity σ, then we can estimate the current source density by measuring Φ at a set of points with extracellular electrodes, and using some method to approximate the second spatial derivative (Łęski et al., 2011; Nicholson and Freeman, 1975; Pettersen et al., 2006). Though progress has been made to improve CSD estimation, it is inherently limited by its characteristic as an inverse problem: an infinite configuration of current sources could result in the potentials measured with a finite number of electrodes (Pettersen et al., 2012). It is therefore not possible to ascribe estimated current sources to currents in specific neurons, only to estimate a value for the average of all current sources at the resolution of the measuring electrode array.

2.2.3 Forward model of extracellular field potentials

An alternative approach to solving the inverse problem described above is to instead specify a model of individual neuron dynamics and then calculate the extracellular potential based on the current sources in this model. This is known as a forward

modelling approach: rather than inferring a model from the data, we instead predict data from a specified model (Einevoll et al., 2013). The advantage of this approach is that the theory of current flow in single neurons is well developed, having been studied extensively since the 1940s (Koch, 1999; Johnston and Wu, 1995). Given that we have a well-established theory to calculate the membrane currents in a neuron of arbitrary shape and size, we can combine this theory with Poisson's equation to calculate extracellular potentials for a given neuron model. This approach has been used to improve CSD estimates in certain cases by constraining the CSD solution with prior knowledge of neuronal morphologies and synaptic connectivity distributions (Einevoll et al., 2007; Gratiy et al., 2011; Pettersen et al., 2008).

To find the extracellular potential, we integrate Poisson's equation (Plonsey and Barr, 2000):

$$\Phi\left(x', y', z'\right) = \frac{1}{4\pi\sigma} \int \frac{I_v\left(x, y, z\right)}{r\left(x, y, z, x', y', z'\right)} dxdydz, \tag{2.3}$$

where x', y' and z' are the coordinates at which the potential is measured, and r is the radius between the potential measurement and the sources (Plonsey and Barr, 2000). If we assume that the current sources are a finite number K of point sources - one per neuronal membrane segment - then equation 2.3 becomes

$$\Phi\left(\mathbf{r}'\right) = \frac{1}{4\pi\sigma} \sum_{k=1}^{K} \frac{I_k}{|\mathbf{r}' - \mathbf{r}_k|}, \tag{2.4}$$

(Pettersen et al., 2012), where we have replaced x', y' and z' with the coordinates vector $\mathbf{r}'$, r with the coordinates $\mathbf{r}_k$ of each point source k, and I_k is the membrane current of point source k. Kirchoff's current conservation law requires that the total current across a neuron's membrane is zero: currents flowing into a neuron are balanced by currents flowing out of it. This requirement leads to the conclusion that to generate an extracellular potential, a model must include neurons with a spatial structure rather than the single compartment neurons frequently used in network simulations (Pettersen et al., 2012).

Current flow in a neuron with a spatial structure can be calculated using the cable equation (Niebur, 2008; Thomson, 1854):

$$\hat{C_m}\hat{R_m}\frac{\partial V}{\partial t} = E_l - V + \frac{\hat{R_m}}{\hat{R_a}}\frac{\partial^2 V}{\partial x^2}, \tag{2.5}$$

where $\hat{C_m}$ is the neuronal membrane capacitance per unit length, $\hat{R_m}$ is the membrane resistance per unit length, V is the membrane potential, E_l is the leakage potential (the membrane potential at which there is no membrane current flow) and $\hat{R_a}$ is the axial resistance per unit length (the resistance along the neural processes). The cable equation assumes that a neuron can be treated as a conducting cable, surrounded by an insulating material (its cell membrane), placed in conducting medium (the extracellular fluid) (Johnston and Wu, 1995; Niebur, 2008). The arrangement of two conductors separated by a thin insulator causes the membrane to act as a capacitor. However, the membrane has a finite resistance, so current can flow across it. The cable equation also assumes that current only flows axially in one dimension: neural processes are long and thin, so the variation in V will be much greater along the long (x) axis than in any other direction (Niebur, 2008).

Equation 2.5 describes current flow in a passive cable. However, neuronal membranes also contain many different ion channels, such as the sodium and potassium channels involved in action potential generation. Localised membrane conductance changes are also effected by synapses positioned over the neuron's dendritic tree. These influences can be inforporated into the equation by adding suitable terms to the right hand side of equation 2.5.

Solving the cable equation for situations more complex than the simple passive cable usually requires the use of numerical methods (Bower and Beeman, 1995; Carnevale and Hines, 2006; Koch and Segev, 1998). We can discretise the equation in space by considering finite-sized lengths Δx of cable, resulting in

$$\hat{C_m}\Delta x\frac{dV}{dt} = -\hat{g_l}\Delta x\left(V - E_l\right) + \frac{V\left(x-\Delta x\right) - V\left(x\right)}{\hat{R_a}\Delta x} - \frac{V\left(x\right) - V\left(x+\Delta x\right)}{\hat{R_a}\Delta x}, \quad (2.6)$$

(Niebur, 2008), where $\hat{g_l} = 1/\hat{R_m}$ is the membrane conductance per unit length. A discrete cable model is called a compartmental model: each compartment is considered to be small enough to have constant membrane potential along its length. The length of each compartment is not necessarily the same - different lengths for different neurite segments can be used. Including terms for active ion currents I_{ion} (for example, action potential generating sodium and potassium currents), synaptic currents I_{syn} and an arbitrary number of adjacent compartments K to take into account branching cables, the equation for the membrane potential of compartment

j is then

$$C_{m,j}\frac{dv_j}{dt} = -g_{l,j}\left(v_j - E_l\right) - \sum_{k=1}^{K} g_{jk}\left(v_j - v_k\right) + I_{ion,j} + I_{syn,j}, \tag{2.7}$$

where $C_{m,j}$ is compartment j's membrane capacitance, v_j is its membrane potential, $g_{l,j}$ is its leak conductance, g_{jk} is the axial conductance between compartments j and k, and v_k is the membrane potential of compartment k. We know from Kirchoff's law that the sum of currents entering and exiting the compartment must be zero. Current flows into/out of the compartment either across its membrane or from its adjacent compartments; therefore, the sum of the axial currents $I_{ax,j}$ at compartment j must be equal to the negative of the sum of the membrane currents $I_{mem,j}$:

$$I_{mem,j} = -I_{ax,j} = \sum_{k=1}^{K} g_{jk}\left(v_j - v_k\right) = -C_{m,j}\frac{dv_j}{dt} - g_{l,j}\left(v_j - E_l\right) + I_{ion,j} + I_{syn,j} \tag{2.8}$$

(Johnston and Wu, 1995). We now have an equation for calculating the membrane current at every compartment in a compartmental neuron model. Finally, we note that a discretisation of the dendrites into compartments as described above results in a tree of connected cylinders. We could approximate the membrane current at each cylinder as a point current-source for the purposes of calculating the extracellular potential, but a more accurate approach is given by treating the cylinders as line-sources (Holt and Koch, 1999; Holt, 1998). The potential due to a line current-source j is given by

$$\Phi_j = \frac{I_{mem,j}}{4\pi\sigma\Delta s_j}\log\left|\frac{\sqrt{h_j^2 + \rho_j^2} - h_j}{\sqrt{l_j^2 + \rho_j^2} - l_j}\right|, \tag{2.9}$$

where Δs_j is the length of the compartment, h_j is the longitudinal distance of the measurement point from the compartment's end point, $l_j = \Delta s_j + h_j$ is the longitudanal distance of the measurement point from the compartment's start point, and ρ_j is the perpendicular distance from the compartment (Holt and Koch, 1999; Pettersen and Einevoll, 2008). We can now use equation 2.8 to calculate the membrane current at each compartment in a compartmental neuron model, and equation 2.9 to calculate the extracellular potential at an arbitrary location due to these currents.

2.2.4 Model assumptions

The forward model described above contains several important assumptions, listed in (Pettersen et al., 2012). We reproduce this list here and offer justifications and implications of each assumption:

1. The quasistatic approximation of Maxwell's equations is valid. The quasistatic approximation assumes that, at the time scales of interest to the neurophysiologist (frequencies up to about 1000Hz), the magnetic and electric fields can be treated as if they were uncoupled. Experimental parameters in biological tissue suggest that this approximation is justified (Plonsey and Heppner, 1967; Rosenfalck, 1969).

2. The extracellular medium is linear. This means that the electric field is linearly related to the current density by the extracellular conductivity.

3. The extracellular medium is resistive. This means that the tissue's capacitive effects are negligible compared with resistive effects.

4. The extracellular conductivity is homogeneous and isotropic. This means that the tissue conductivity is the same at all points in space and in all directions. Experimental measurements suggest these are reasonable assumptions in cortical grey matter, though not in white matter (Logothetis et al., 2007). As we are investigating local cortical circuits, we consider this assumption to be reasonable. Recent results have opened up discussion on points 3 and 4: current-source density methods appeared to reveal the presence of current monopoles in rat barrel cortex (Riera et al., 2012). While this result could be explained by methodological artefacts as current monopoles are not possible in resistive homogeneous biological tissue (Gratiy et al., 2013), monopoles could be generated if assumptions 3 or 4 are false (Bedard and Destexhe, 2013).

5. The extracellular conductivity is not frequency-dependent. This assumption is currently under discussion in the literature. Experimental studies have found conflicting evidence regarding the frequency dependence of brain tissue. Bédard et al. have used various methods to infer a low-pass filtering effect inherent to the neural tissue itself (Bédard et al., 2010; Dehghani et al., 2010). Other experiments, including direct measurements in macaque brain tissue,

have suggested that at frequencies of relevance to neurophysiology any tissue filtering effects are negligible (Logothetis et al., 2007). Extracellular recordings do show low-pass characteristics; while this could be due to tissue filtering (Bédard et al., 2006b, 2004, 2010), an alternative explanation is that dendritic filtering of the membrane current causes the low-pass filtering effect on the extracellular potentials (Łęski et al., 2013; Lindén et al., 2010; Pettersen and Einevoll, 2008).

As well as these assumptions, we make some further simplifications when simulating extracellular potentials. We are particularly interested in the low frequency part of the extracellular signal: the local field potential (LFP). It is usually assumed that the membrane currents contributing to the LFP are primarily due to synaptic activity, as the main active conductances in the membrane have fast dynamics (e.g. spike generation). Our simulations therefore only considered passive compartmental models with synaptic input, i.e. we set I_{ion} in equations 2.7 and 2.8 to zero. A recent modelling study has suggested that, in fact, active membrane currents can contribute prominently to the LFP at frequencies as low as 50Hz (Reimann et al., 2013). This is an important observation given the general assumption that LFPs are dominated by synaptic activity. However, the results should be interpreted with some caution. In a large and complex model as used in this study, very many different parameter sets can give rise to very similar activity (Prinz et al., 2004). The ion channel parameters of the individual neuron models were constrained using intracellular rather than extracellular measurements, and while care was taken to constrain these using measurements in both soma and dendrites, the parameter ranges that produced the correct membrane potential responses were quite broad (Hay et al., 2011). Synaptic parameters were chosen apparently arbitrarily, so might not have reflected the true synaptic strengths in the kind of network regimes that were simulated. A more systematic approach is required to truly pin down the contribution of active currents. Another barrier to such investigations is the computational cost of simulating realistic active compartmental models: the simulations in (Reimann et al., 2013) were performed on a supercomputer with 4 096 processors, with a second of simulation time requiring 45 minutes to compute. Models considering only synaptic activity and passive return currents are still very useful for interpreting experimental results: interpretations of experimentally recorded LFPs usually assume that the signal is

mostly generated by these currents, so a model of the experimental setup that incorporates only these currents can reveal to what extent this assumption is true for that particular experiment.

2.3 Simplified neuron models for network simulations

2.3.1 Need for a simplified model

Previous simulation work has used reconstructed neuron models each containing hundreds of compartments to model LFPs (Einevoll et al., 2007; Gold et al., 2006; Gratiy et al., 2011; Holt and Koch, 1999; Łęski et al., 2013; Lindén et al., 2010, 2011; Pettersen and Einevoll, 2008; Reimann et al., 2013). Network models of such reconstructions including active conductances require massive computational resources to simulate, as described above. Even with simple passive dynamics, the large number of compartments in each model makes simulating networks of reconstructed neurons infeasible with the resources available to most labs. We wanted to be able to run simulations on commonly available hardware to enable a wide range of theoretical predictions to be made and to allow experimentalists to work directly with the models on their own computers. We therefore needed to choose a suitable method to reduce the recostructed compartmental models to fewer compartments while preserving the crucial features of the LFP that are due to the neuron shape and size.

2.3.2 Considerations for model reduction

As we intend to use the models to investigate the LFP across space, the reduced model should preserve the spatial features of the LFP generated by synaptic input on an individual neuron. This means that the currents flowing out of the neuron's dendritic tree in response to an input in the reduced model should have a similar spatial profile and magnitude to those currents in the full model. This will depend on the reduced model's spatial extent being similar to the full model, and the degree to which the membrane resistance and capacitance, as well as the axial resistance, are conserved.

We chose to investigate the compartmental reduction method of (Bush and Sejnowski, 1993). This method relies on conserving axial resistance, collapsing multiple compartments into a single compartment based on their electrotonic distance from the cell soma. The reduced compartment's length is equal to the mean length of the compartments it represents in the full model, so that the reduced model is of a similar spatial size to the full model. Membrane resistance and capacitance are not conserved, as the total membrane area is reduced. However, these values can be scaled appropriately by multiplying the capacitance and dividing the membrane resistance of each compartment by the ratio of the full model's total membrane area to the reduced membrane area (Bush and Sejnowski, 1993). Current inputs at various locations on both the full and reduced models show similar voltage attenuation at the soma, indicating that the current flow in the equivalent reduced model is similar to that in the full model. We therefore decided to test how well the spatial features of the LFP were preserved when performing equivalent simulations with the full and reduced models. While the reduced number and lateral spread of the compartments will alter the LFP due to identical synaptic input on a single neuron level, we anticipated that the overall LFP spatial characteristics from a population would be similar.

During our investigations with the reduced cell models, we noticed an inconsistency between the reported compartment sizes in (Bush and Sejnowski, 1993) and the reported reduction method. The soma compartment of each reduced model was supposed to be the same size as the cell reconstruction on which it was based, but the reduced model soma dimensions reported in (Bush and Sejnowski, 1993) do not match with the original models. Additionally, a reduction of a layer 4 spiny stellate cell was not created in (Bush and Sejnowski, 1993) . We therefore recalculated the compartment lengths and diameters from the three cell types specified in (Mainen and Sejnowski, 1996) using the method specified in (Bush and Sejnowski, 1993). For these calculations, we used a version of the NEURON code originally written by Alain Destexhe to reduce a compartmental model to 3 compartments (Destexhe, 2001), modified by Michael Hines to work for any number of compartments. Michael Hines' version of this code is available on the NEURON forum at http://www.neuron.yale.edu/phpbb/viewtopic.php?f=13&t=589 (last accessed 16th January 2014). The recalculated cell dimensions are given in Table

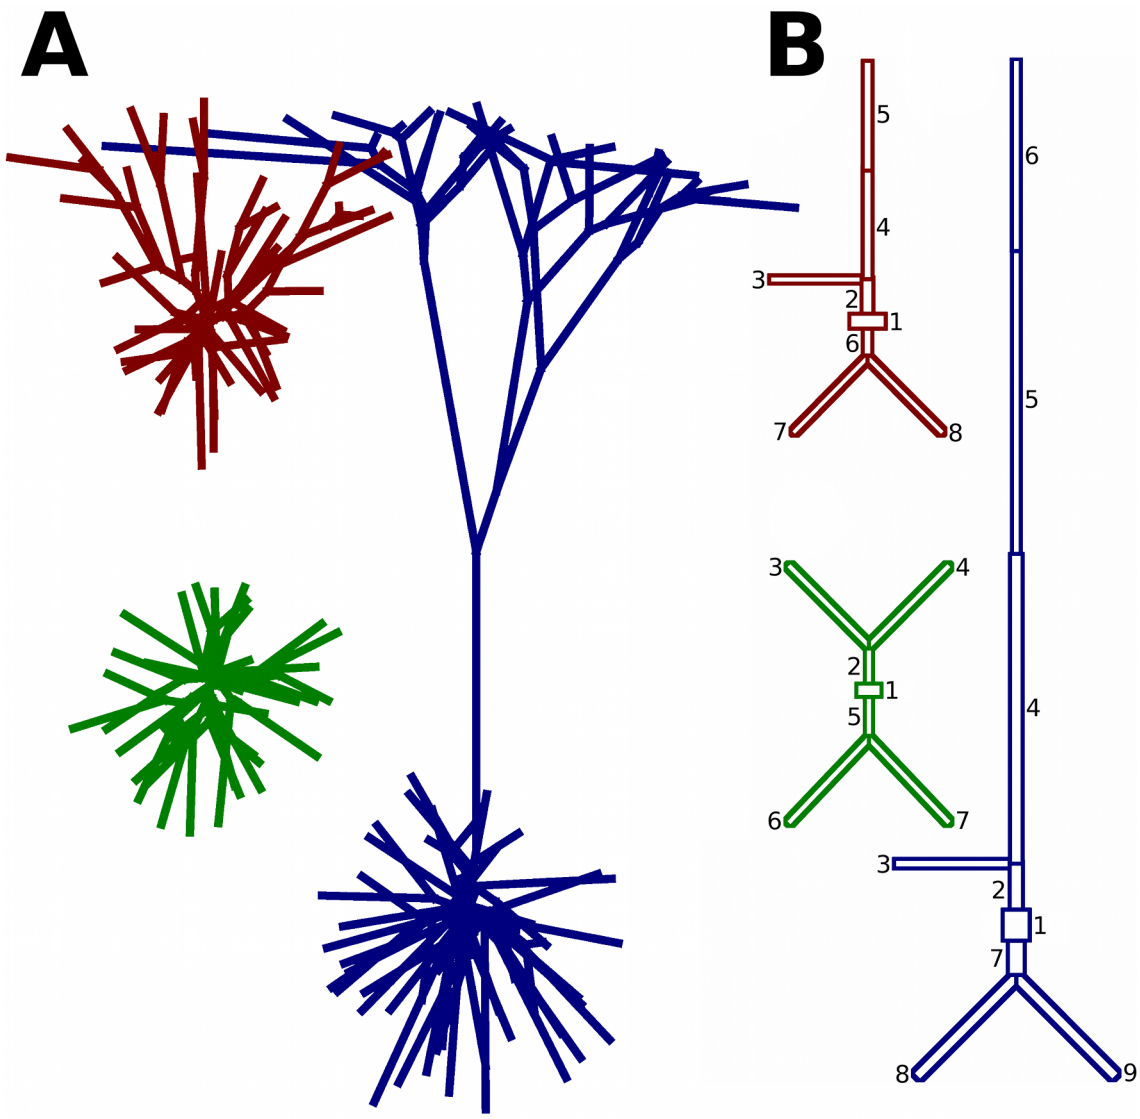

Figure 2.1: (A) Comparmental structures of morphological cell reconstructions from (Mainen and Sejnowski, 1996). The layer 2/3 pyramidal cell, layer 4 spiny stellate cell and layer 5 pyramidal cell are shown in red, green and blue, respectively. (B) Compartmental models reduced from the structures in (A) according to the method in (Bush and Sejnowski, 1993). Compartment numbers correspond to those in table 2.1.

	Layer 2/3 pyramidal		Layer 4 spiny stellate		Layer 5 pyramidal	
Compartment number	Length	Diameter	Length	Diameter	Length	Diameter
1	13	29.80	10	24.00	35	25.00
2	48	3.75	56	1.93	65	4.36
3	124	1.91	151	1.95	152	2.65
4	145	2.81	151	1.95	398	4.10
5	137	2.69	56	1.93	402	2.25
6	40	2.62	151	1.95	252	2.40
7	143	1.69	151	1.95	52	5.94
8	143	1.69	-	-	186	3.45
9	-	-	-	-	186	3.45

Table 2.1: Compartment dimensions of the reduced model neurons, calculated using the method from (Bush and Sejnowski, 1993). Note: these values are different from the values given in (Bush and Sejnowski, 1993), as we noticed discrepancies in the originally reported values so re-calculated the dimensions. Membrane resistance/-capacitance scale factors (see main text) were 2.96 for the L2/3 pyramidal neuron, 2.93 for the L4 spiny stellate neuron, and 2.95 for the L5 pyramidal neuron.

2.1. The membrane resistance/capacitance scale factors for the re-calculated reduced morphologies were 2.96 for the L2/3 pyramidal neuron, 2.93 for the L4 spiny stellate neuron, and 2.95 for the L5 pyramidal neuron.

2.4 Validating the simplified LFP generation model

We tested the effects of this reduction on the generated LFP by reproducing the experiments detailed in (Lindén et al., 2011). Ten thousand model neurons with the same morphology were positioned randomly within a 1 mm radius cylinder, with uniform spatial distribution and constant soma depth. The soma depth is chosen as the centre of the layer in which the neuron's soma resides, with layer boundaries given in (Stepanyants et al., 2008). One thousand synapses (excitatory, current-based, single exponential type with time constant 2 ms and fixed amplitude 50 pA) were placed randomly on the dendritic compartments of each neuron, with uniform density with respect to membrane area. Each synapse received an independent Poisson spike input train with a rate of 5 Hz to simulate uncorrelated synaptic input. LFPs were calculated at the centre of the population, at five depths corresponding to the centre of each layer. Each neuron's LFP contribution was calculated and stored individually, so that the compound LFP could be calculated for different population

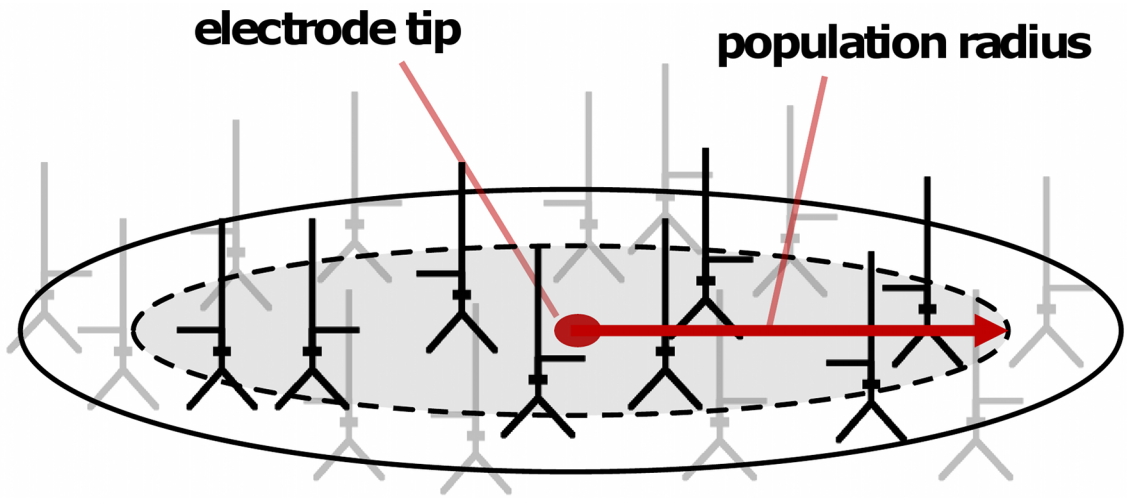

Figure 2.2: Measuring the spatial reach of the LFP. A population of neurons receiving synaptic input is positioned within in a circle surrounding a central point (shaded region), and the LFP due to this population is calculated at that point. We add neurons to increase the population radius up to a maximum of 1 mm (solid black circle), measuring the LFP with every 0.05 mm radius increase. The radius at which the LFP magnitude reaches 95% of its value at the maximum population radius is then defined as the LFP spatial range: neurons further away from the electrode than this distance only contribute minimally to the LFP. This provides an "electrode-centric" definition of LFP range.

radii, up to the maximum population radius of 1 mm. The magnitude of an LFP signal was defined as its standard deviation. The LFP range at each depth was calculated by varying the population radius from 0 mm to 1 mm in 0.05 mm steps, and measuring the radius at which the LFP magnitude at that depth reached 95% of its value at the maximum 1 mm radius (Lindén et al., 2011) - see Figure 2.2.

We repeated this procedure for the three neuron types used by Lindén et al.: layer 2/3 (L2/3) pyramidal, layer 4 (L4) spiny stellate, and layer 5 (L5) pyramidal. To minimise variability due to random positioning, we ensured that the x- and y-coordinates for the position of each neuron's soma were the same for each population. We also added the same synapses to the neurons at the same coordinates in the different populations, to minimise variability in the synaptic input distribution between conditions. We compared LFPs generated by the morphological reconstructions of these neuron types from (Mainen and Sejnowski, 1996) - hereafter referred to as Mainen cells - with the LFPs from reduced versions of these models created using Bush and Sejnowski's method (Bush and Sejnowski, 1993) - hereafter referred to as Bush cells. We also compared results with morphological reconstructions of other L2/3 pyramidal and L5 pyramidal cells from different labs, as a crude check of

how the inaccuracies introduced by the model reduction compared with differences due to differing morphology. The extra comparison morphologies used were downloaded from the NeuroMorpho.org database (Ascoli et al., 2007). We used a cat visual cortex layer 2/3 pyramidal neuron from (Kisvárday and Eysel, 1992) (NeuroMorpho ID NMO_00856) - hereafter referred to as Kisvárday cells - and a cat visual cortex layer 5 pyramidal neuron from (Contreras et al., 1997) (NeuroMorpho ID NMO_00880) - hereafter referred to as Contreras cells. We used the following passive parameters for each simulation: specific axial resistance $R_A = 200\,\Omega \cdot \mathrm{m}$, specific membrane resistance $R_M = 20\,\mathrm{k}\Omega \cdot \mathrm{m}^2$, and specific membrane capacitance $C_M = 1\,\mu\mathrm{F} \cdot \mathrm{m}^{-2}$. The Bush neurons had their specific membrane resistance and capacitance scaled by the relevant scale factors for each cell type as described above.

We implemented these experiments using LFPy (Lindén et al., 2014), a tool for simulating extracellular potentials from single neurons written in Python. LFPy uses the NEURON simulator (Hines and Carnevale, 1997) as a back-end for calculating membrane currents at each compartment of a neuron model. Before conducting the experiments, we slightly modified LFPy to control how NEURON dealt with the compartmentalisation of the neurons. NEURON's implementation of compartmental modelling attempts to separate the biological issue of morphology from the numerical issue of spatial discretisation (Carnevale and Hines, 2006). Users therefore specify the morphology in terms of connected "sections", which are then broken down into one or more "segments", which is how NEURON refers to compartments. NEURON automatically calculates the number of segments to break each segment into based upon certain accuracy criteria, or into an arbitrary number specified by the user. As the definition of the reduced cell models is in terms of a specific number of compartments rather than a dendritic branching morphology, we wanted to make sure that NEURON would use one segment per section in the Bush cell model definition. The LFPy interface does not let the user specify an arbitrary number of segments per section, so we added this functionality and ensured that we ran all our experiments with one segment per section for both the Mainen and Bush cells. Our initial results obtained for the Mainen cells were similar to those reported in (Lindén et al., 2011), suggesting that the inaccuracies introduced by our coarser spatial discretisation were minimal.

2.4.1 LFP range and magnitude

The results of these experiments are shown in Figure 2.3. For each neuron type, the LFP range and magnitude in each layer for the population of Bush cells are close to those for the population of Mainen cells. The LFP range is smallest in the soma layer (<250 microns) with the range increasing in the layers above and below the soma, while the LFP magnitude is largest in the soma layer and decreases in the layers above and below the soma. The differences between the results for the L4 spiny stellate models is small, so we concentrate on the pyramidal neuron population results.

For the L2/3 pyramidal neurons, the LFP spatial range in the soma layer is very similar between the Bush and Mainen populations, but above and below this layer the discrepancy increases, with the largest difference of ∼200 microns in L1. The range differences in all other layers are <100 microns. For the L5 pyramidal neurons, the LFP spatial range difference is again smallest in the soma layer, and approximately 150 microns in layers 4, 2/3 and 1. These differences are comparable to the differences in LFP range between the Mainen population and the Contreras population in L5, and the Mainen population and Kisvárday population in L2/3. The differences in LFP magnitude in each layer are also comparable. This suggests that, while the reduced models are not ideal substitutes for the morphological reconstructions, the errors incurred by the reduction method are similar to those introduced by neglecting morphological diversity in reconstructed neuron model populations. The general profile of the LFP across the layers, at least, is preserved adequately.

The results above were generated using uncorrelated synaptic inputs over the entire dendritic tree of each neuron in each population. However, we also needed to check the situation for highly correlated, spatially localised synaptic inputs. We are ultimately interested in using these compartmental reductions in network models of gamma oscillations, in which pyramidal neurons receive highly correlated inhibitory synaptic input to their perisomatic region. We therefore repeated the previously described experiments measuring the LFP magnitude and range, but positioned each neuron's 1000 synapses onto its soma compartment (we only repeated the simulations for the pyramidal neuron morphologies, as the LFP spatial profile for the spiny stellate cells was shown not to change significantly with correlated input (Lindén et al., 2011)). In the previous experiments with no correlations between

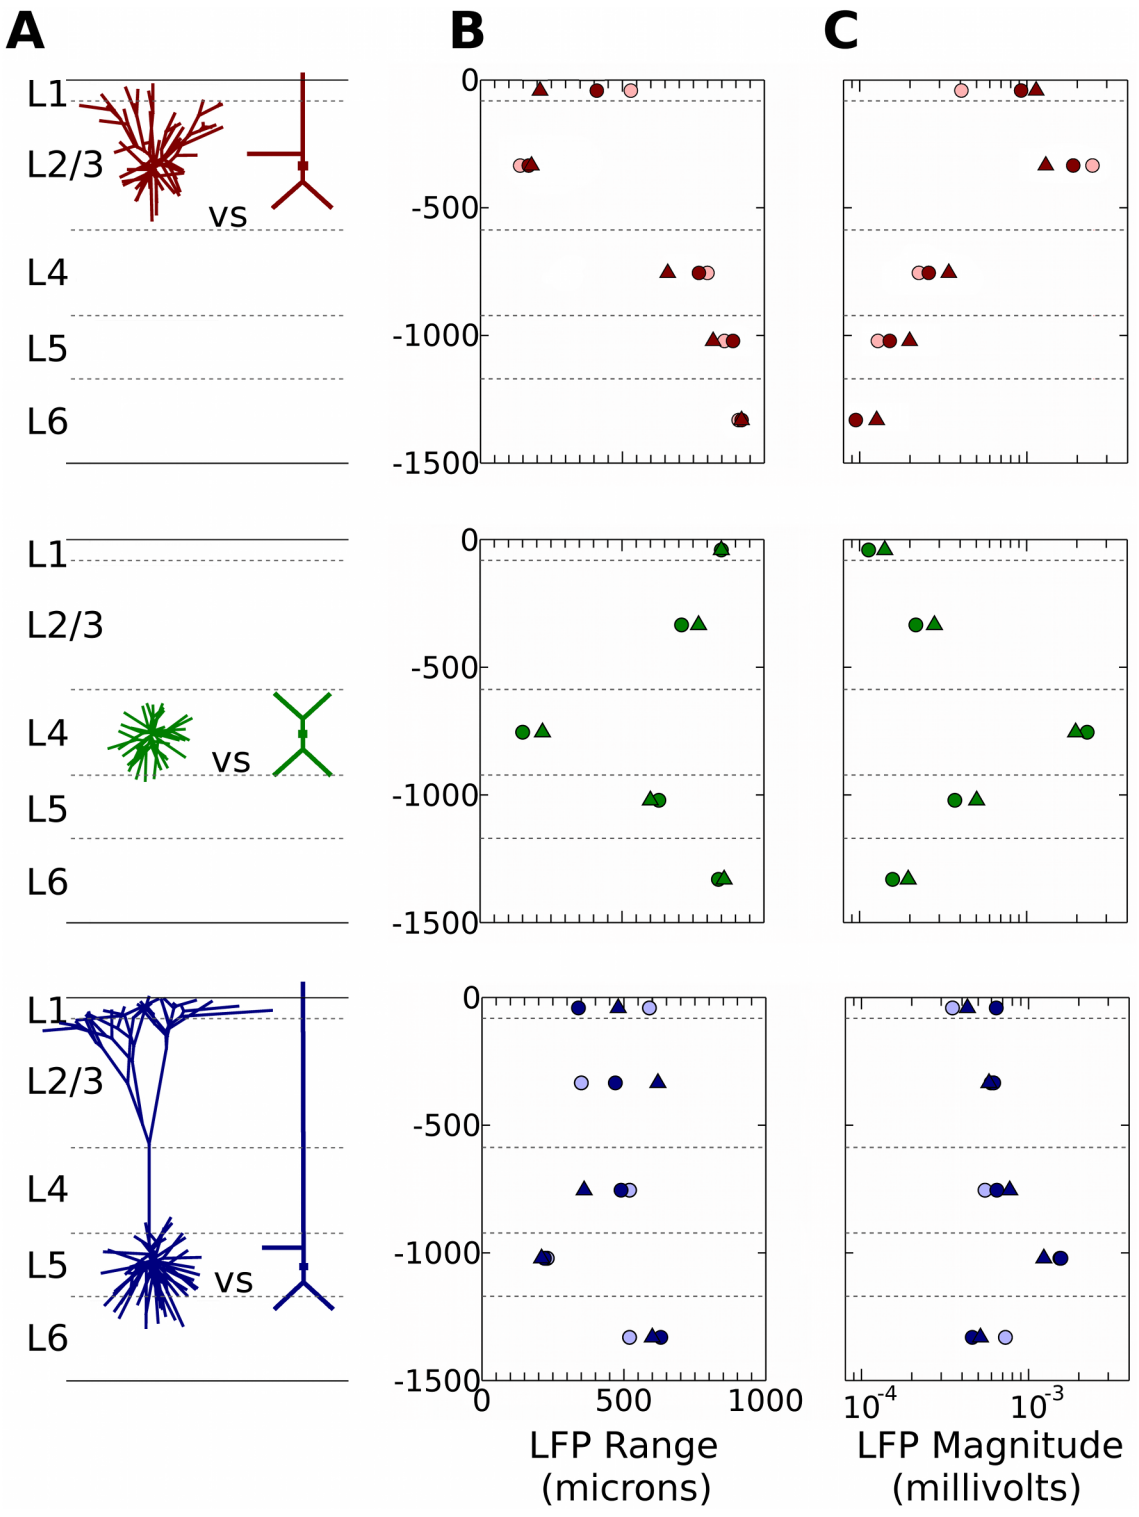

Figure 2.3: Comparison of simulated LFPs from the Bush and Mainen cell models. Top (red): L2/3 pyramidal neuron, middle (green): spiny stellate cell (morphology also used for interneurons), bottom (blue): L5 pyramidal neuron. (A) Comparison of original and reduced multi-compartment models of each neuron type. (B) Range and (C) magnitude of simulated LFPs. Circles show values for the original cell reconstruction populations, triangles for the reduced neuron model populations. Light blue circles in bottom panel show values for a different reconstructed L5 cat pyramidal neuron. Grey dashed lines show layer boundaries. All y-axis values in microns.

synaptic inputs, each synapse was assigned an independent Poisson spike train, for a total of $10000 \times 1000 = 10^7$ independent spike trains at 10^7 synapse locations. To introduce correlations in the synaptic input, we followed the method described by Lindén et al. (2011) and Łęski et al. (2013). Each synapse in the model was now assigned a spike train drawn without replacement from a finite pool of pre-generated spike trains. By reducing the number of Poisson spike trains in the pool so that some synapses shared a common input pattern, we could control the level of input synchrony to the neurons. The resulting input correlation is given by the total number of synapses per neuron divided by the number of independent spike trains (Łęski et al., 2013). To simulate highly correlated input, we used 2000 independent spike trains, resulting in an input correlation of $1000 \div 2000 = 0.5$ (in other words, any two neurons share on average $1000 \times 0.5 = 500$ common input spike trains).

For these simulations, we also introduced random variability in the soma depth of the neurons. In real cortical tissue, neuronal cell bodies are not all at the same depth within a layer, but are spread out between the layer boundaries (it would be impossible for them to be at the same depth given usual measures for neuronal density and soma diameter). A previous forward modelling study used a gaussian depth profile, with the maximum cell body density in the centre of the layer, which decayed towards the layer boundaries (Gratiy et al., 2011). However, it is unclear whether this distribution is realistic, so we simply chose a uniform depth distribution as this appears qualitatively to fit the distribution of cell bodies in Golgi-stained cortical slices (Abeles, 1991). We distributed L2/3 pyramidal neuron somas between -334 microns and -534 microns, and L5 pyramidal neuron somas between -970 microns and -1170 microns from the cortical surface. These ranges ensured that the neuron somas remained within the correct layer boundaries, and that their apical dendrites were not positioned above the cortical surface.

Figure 2.4 shows the spatial profiles of LFP for the different populations. In these simualations, we measured the LFP at 50 micron intervals, to see how well the Bush models preserved the LFP at this level of detail. We used 12 electrode points in L1 and L2/3 with the L2/3 populations, and 26 electrode points spanning all layers with the L5 populations. Both the range and magnitude profiles show that the LFP from the Bush population matched the LFP from the Mainen population well, again within the bounds the LFP profile of Kisvarday and Destexhe comparison

populations. The minimum range and magnitude in the L2/3 populations are just above the minimum soma depth, and a few hundred microns above the minimum soma depth in the L5 population. This depth is where the synaptic currents at the soma are approximately balanced by the opposite return currents in the dendrites; below and above this minimum point, the somatic or the apical dendritic currents dominate the LFP signal, respectively.

2.4.2 LFP frequency scaling

We also checked the power spectra of the simulated LFPs to make sure the Bush model populations reproduced similar frequency scaling properties to the Mainen cell populations. Results for the case of uncorrelated synaptic input given in Figure 2.5 show that, in each layer, the 95% confidence intervals for each model type overlap over the range of frequencies from 2-450 Hz (the overlap continues down to 1 Hz; this is not shown in order to improve the plot resolution at higher frequencies).

Results with correlated inputs at the soma compartment again show substantial overlap of the 95% confidence interals for the power spectra at each electrode. Figure 2.6 A shows comparison at each electrode used in L2/3, while Figure 2.6 B only shows results at every second electrode for the L5 simulations (100 micron spacing), for ease of visualisation. The biggest discrepancy between the LFP power spectra for each model occurs around the level of the LFP range minimum. The LFP power up to 100 Hz is reliably reproduced at every measurement point, and up to 450 Hz at all but one point with the L2/3 populations. This point corresponds to the point at which the LFP range and magnitude are lowest. The reduced accuracy at higher frequencies in the L5 models should be taken into account if frequencies above 100 Hz are analysed in models containing L5 pyramidal cells.

2.4.3 When is the simplified model not suitable?

Our results show that, despite the crude approximation using <10 compartments, the reduction method proposed in (Bush and Sejnowski, 1993) creates neuron models that reproduce the population LFP characteristics of the full neuron morphologies on which they are based. Discrepancies in simulated LFPs between the full and reduced model populations are generally small, and of a similar magnitude to differences

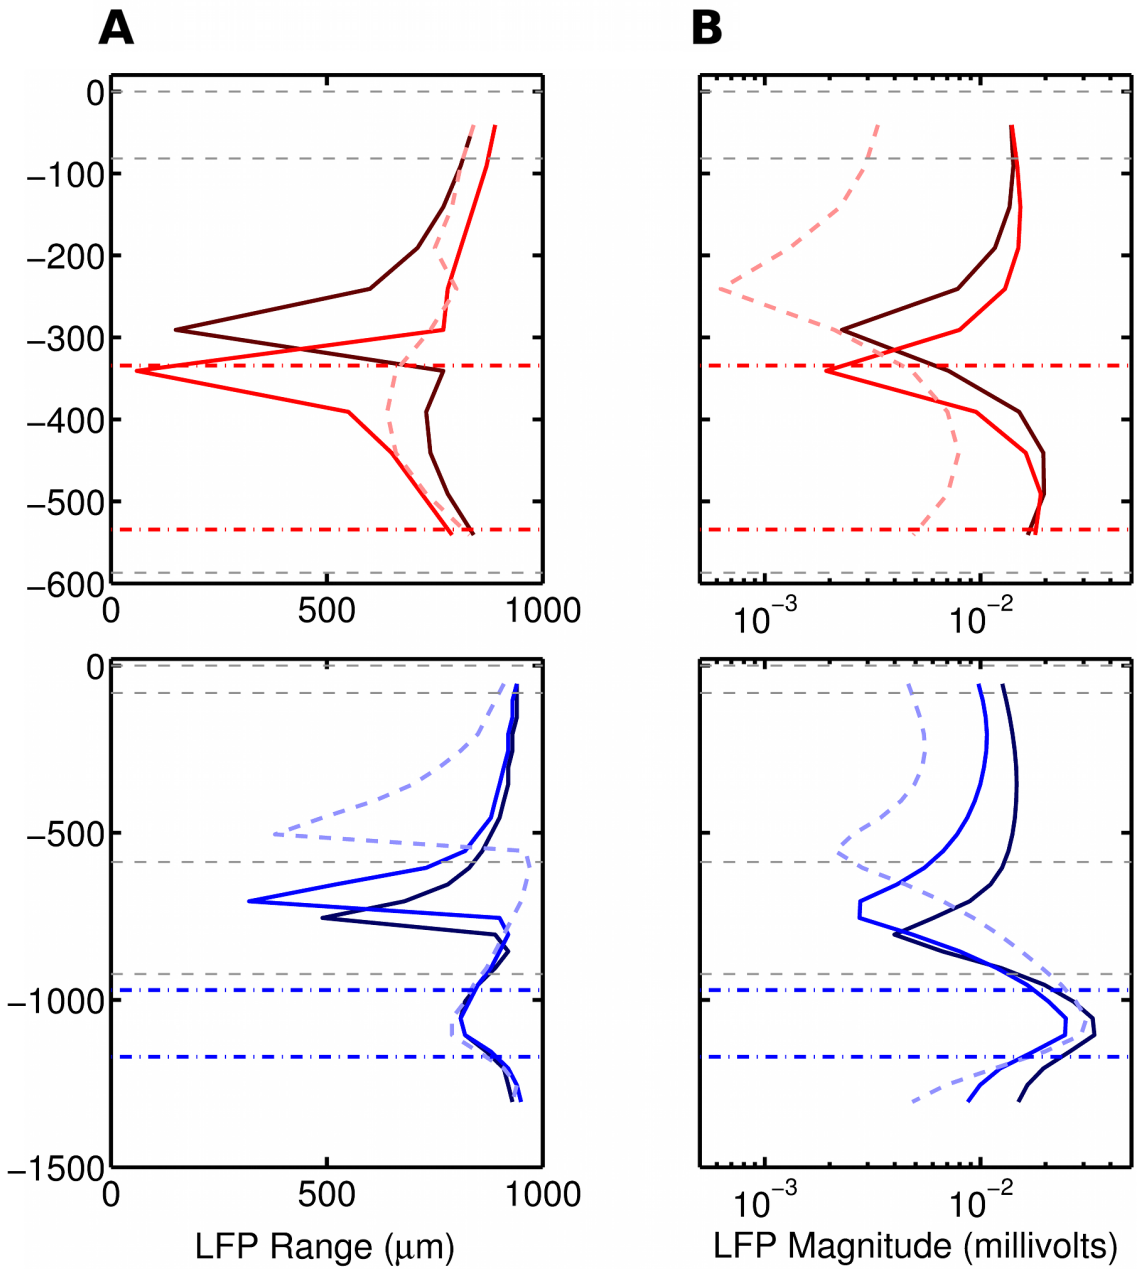

Figure 2.4: Comparison of simulated LFPs from the Bush and Mainen cell models for highly correlated input at the soma compartment. Top (red): L2/3 pyramidal neuron, bottom (blue): L5 pyramidal neuron. (A) Range and (B) magnitude of simulated LFPs. Bright red/blue lines show range and magnitude values for the Mainen cell populations, dark red/blue lines show range and magnitude values for the Bush cell populations. The faded red/blue dashed lines show these values for the Kisvarday cell population in L2/3 and the Destexhe cell population in L5. Grey dashed lines show layer boundaries, red/blue dot-dashed lines show the maximum and minimum soma depths. All y-axis values in microns.

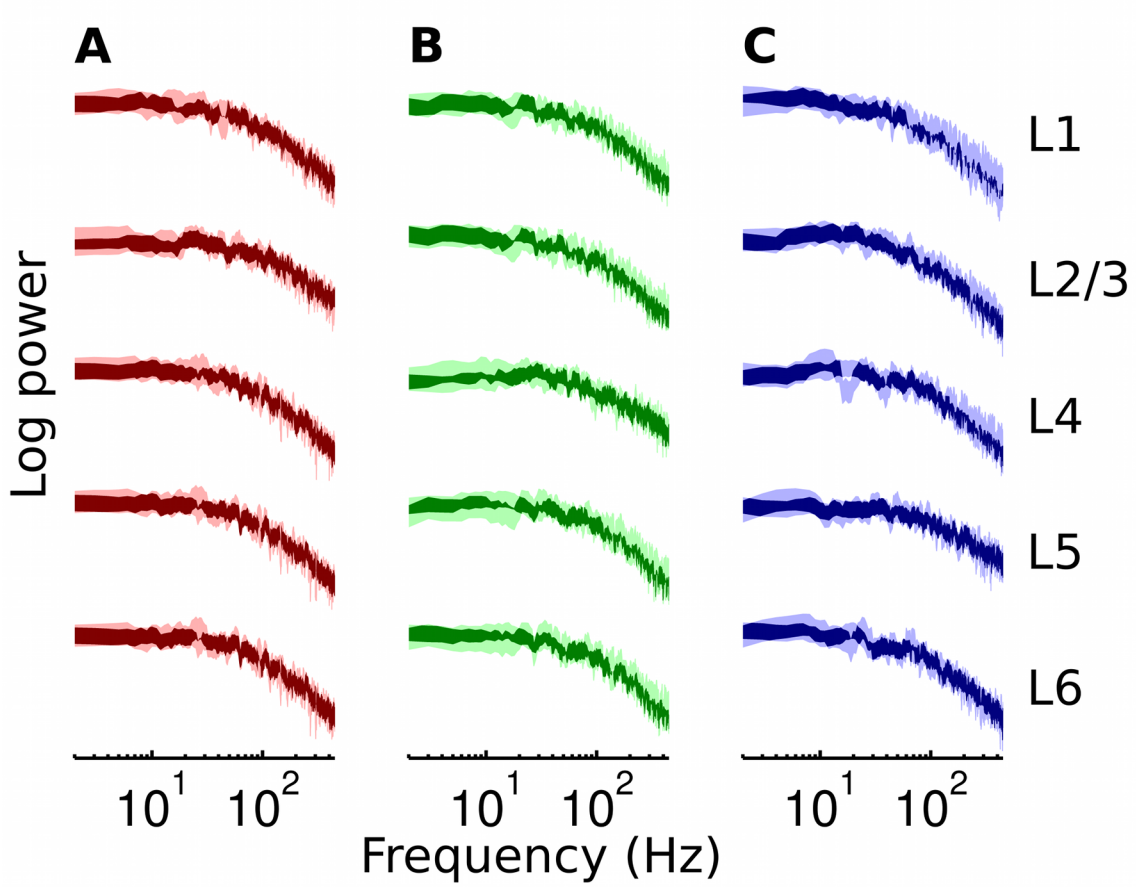

Figure 2.5: Overlap of the 95% confidence intervals for the estimated LFP power spectra produced by each population in each layer shaded dark (uncorrelated input over dendritic compartments). Non-overlapping sections of the 95% confidence intervals are shaded light. Power is plotted in dimensionless, normalised units for ease of comparison. (A) L2/3 pyramidal neuron comparison. (B) L4 spiny stellate cell comparison. (C) L5 pyramidal neuron comparison.

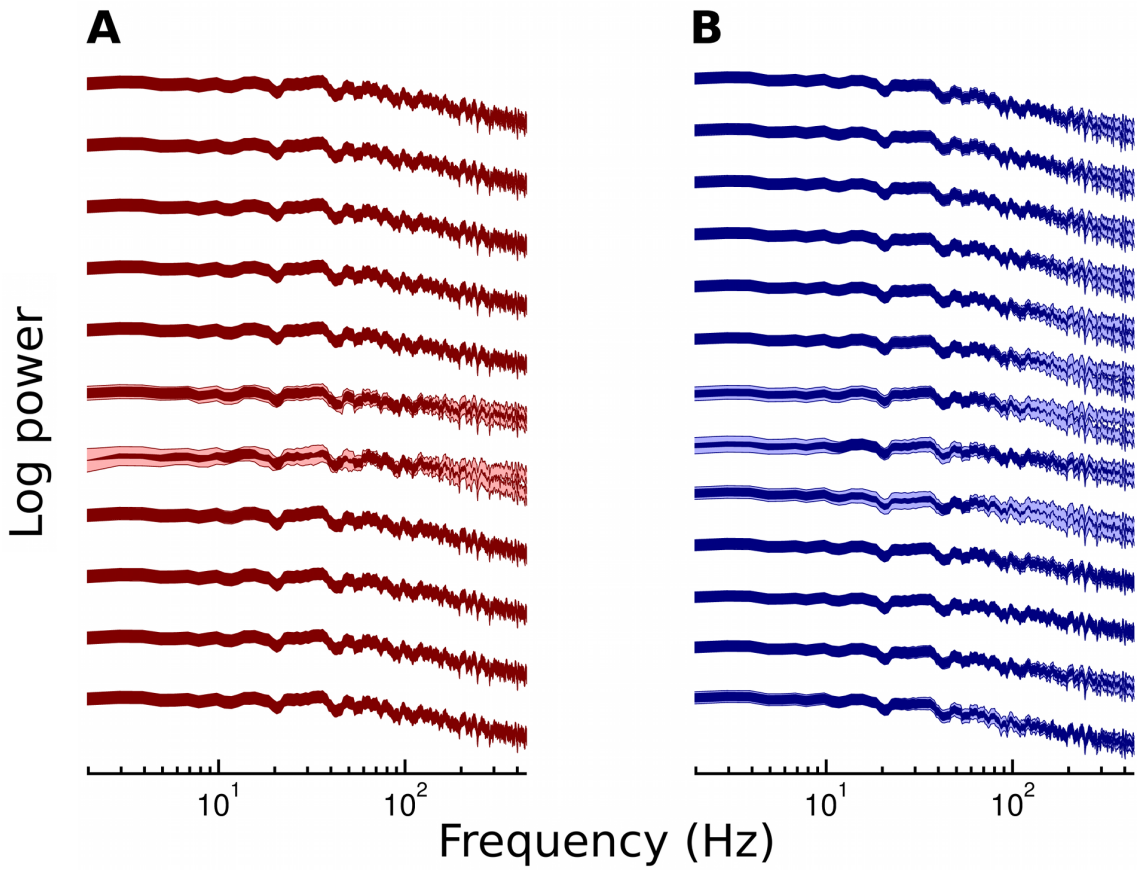

Figure 2.6: Overlap of the 95% confidence intervals for the estimated LFP power spectra produced by the layer 2/3 and layer 5 pyramidal neuron populations at each electrode location shaded dark (correlated input at soma). Non-overlapping sections of the 95% confidence intervals are shaded light. Power is plotted in dimensionless, normalised units for ease of comparison. (A) L2/3 pyramidal neuron comparison. (B) L5 pyramidal neuron comparison (comparisons for 13 out of the 26 LFP measurement points are shown for ease of visualisation).

observed when using models containing different reconstructed morphologies of the same neuron type. This suggested that errors from model reduction were within the bounds of general biological variability.

However, the reduced models would not be suitable in general, and the limitations of the model should be acknowledged when simulating the LFP. We have not tested the full range of possible synaptic inputs, but two extremes: uncorrelated, distributed input and correlated, focal input. It may be that for other input arrangements, inaccuracies introduced by the model reduction would be more apparent. This seems unlikely, though, given the similarity of the LFP between models for two very different input regimes. Secondly, the differences in the LFP have only been measured at the population level. The LFP measured close to a Mainen cell could be quite different to the LFP close to a Bush cell receiving the same input, as overall morphology will have a greater effect on proximal LFP measurements. Investigations into the LFP generated by single neurons, or small groups of neurons, should therefore use morphologically reconstructed cell models. Thirdly, if the frequency scaling profile above 100 Hz is important to the particular investigation, the discrepancies shown in Figure 2.6 should be taken into account. Finally, the reduced models have only been tested with passive dynamics. Different neuron types exhibit different ion channel densities across their dendritic trees, which would affect the LFP depending on the distribution and dynamics of the resulting currents. As the model reduction method does not specify how active conductances would be mapped to the reduced compartments, this would need to be established and the implications for the LFP investigated before the reduced models could be used in such simulations.

2.5 Summary

The extracellular potential is created by neuronal membrane currents. Its distance-dependence means that the signal depends on the position and morphology of the neurons and the position of synapses on the neuron's dendrites relative to the recording point, as well as the activity of the synapses and any active membrane currents. Network models are increasingly used to examine the activity of neuronal networks, but tend to use simplified neuron models that are not suitable for simu-

lating spatially realistic extracellular potentials. As morphologically reconstructed neuron models containing hundreds or thousands of compartments are too computationally expensive to implement in network simulations, we investigated reduced compartmental models with <10 compartments as a compromise between accuracy and efficiency for simulating LFPs resulting from network activity. We found that these reduced models preserved the key spatial and frequency scaling features of the equivalent morphological reconstructions.

While LFPy provides an excellent solution for simulating LFPs from single neurons, it currently does not allow the simulation of LFPs in network models. In the next chapter, we describe the neuronal network simulator, VERTEX, that we developed to simulate LFPs resulting from network activity.

Chapter 3

The VERTEX simulation tool

3.1 Introduction

In the previous chapter, we introduced the theory of extracellular potential generation in the brain, and showed that reduced compartmental neuron models can be used to model the extracellular potential resulting from neuronal network activity. In this chapter, we describe the simulation tool we developed to implement such networks: the Virtual Electrode Recording Tool for EXtracellular potentials (VERTEX). We developed VERTEX to allow us to construct and simulate dynamics in large networks of compartmental neuron models, and to simulate the local field potentials arising from this activity. Even with reduced neuron models, large networks are computationally expensive to simulate, so we used parallel programming and code vectorisation techniques to reduce simulation times. This allowed us to use VERTEX to calculate LFPs in network simulations containing hundreds of thousands of reduced compartmental neurons.

Here we explain our rationale for developing VERTEX, and provide a description of the software, including how networks are initialised, how dynamics are implemented, and the parallel programming and vectorisation methods used. In the next chapter, we describe a large-scale neocortical model that we developed using VERTEX to use in conjunction with experimental multi-electrode array data recorded *in vitro*.

3.2 Need for a new simulation tool

A researcher is faced with many choices when beginning a new modelling project. Models of complex systems not only require mathematical specification, but also simulation, as many of the relevant equations are not solvable analytically. As researchers within the same field frequently address similar problems, various software tools have emerged that facilitate the specification and simulation of commonly used models. Neuroscience is no exception, with hundreds of freely available tools available to assist with creating models of neural systems at every conceivable scale (see, for example, the list at http://www.incf.org/resources/research-tools). An important choice for the researcher, then, is which software tools to choose for their particular modelling purpose, or whether to write custom simulation code from scratch. This choice is influenced by the suitability of current tools for implementing the models of interest, in terms of features, computational speed, software maturity and ease of use.

Chapters 1 and 2 outlined our areas of interest and the approach we decided on taking to investigate them: simulating the electrical dynamics of networks of multi-compartment spiking neurons. Spiking neuronal networks are becoming increasingly important for understanding the behaviour of real neuronal networks as recent advances in computing power enable the simulation of large networks in reasonable time. Many tools already exist for running these kinds of simulations at various levels of detail; for a review see (Brette et al., 2007). Efforts have also been made to allow the specification of a model separately from any particular simulator, meaning that models can be more easily shared between researchers using different simulation tools. While some of these cross-simulator languages are now reasonably mature (Davison et al., 2009; Gleeson et al., 2010), they are still undergoing development and were too incomplete for our purposes when the project started.

In Chapter 2, we assessed the suitability of reduced compartmental neuron models for simulating LFPs using the LFPy tool. LFPy provides Python classes and functions to calculate extracellular potentials at predetermined locations based on the membrane current at each compartment of a NEURON-simulated cell model. This allows the user to leverage the power of NEURON to simulate extracellular potentials from detailed compartmental single cell reconstructions. We showed in

Chapter 2 that, given some caveats, simulation speed can be greatly increased while retaining some important features of interest in the simulated extracellular potentials using reduced compartmental models. Our simulations so far have been using predefined synaptic release times rather than calculating the synaptic behaviour resulting from spiking activity in a connected network. However, we are particularly interested in the LFP resulting from network activity. The NEURON simulator has features for network simulations, as well as parallel simulation to improve speed, but the LFPy interface is not designed to be used in this way (Lindén et al., 2014). An approach taken by Lindén et al. (2011) to calculate extracellular potentials from a specific network was to record incoming spikes at each synapse in a network simulation performed in the NEST simulator using point integrate-and-fire neurons, then to replay these events at synapses in detailed compartmental models implemented with LFPy. However, while NEST can simulate networks of tens of thousands of neurons quickly, the calculations in LFPy take a relatively long time for large numbers of neurons, even if using the reduced compartmental models investigated in the previous chapter. These factors, along with an enthusiasm to learn more about the entire process of simulation and youthful naivety regarding the amount of work involved in building a new simulator, led to the development of VERTEX.

Preliminary simulations performed at the start of the PhD were written in Matlab, and as the project progressed the VERTEX simulation tool evolved from this original code-base. Matlab is not known for its execution speed, so may seem like a strange choice for solving large systems of differential equations as required in neuronal network simulation. However, performance can be dramatically improved through code vectorisation, which minimises the impact of code interpretation overheads. The Matlab Parallel Computing Toolbox allows further performance improvements by providing a very simple way to parallelise computations on multicore computers or over networks. These factors, as well as its ease of use, popularity in the neuroscience community, the ability to perform simulations and analysis in the same environment, and the well-developed interface for integrating C or Fortran functions for future performance enhancements influenced our decision to continue the development of VERTEX in Matlab.

3.3 Description of the VERTEX simulator

3.3.1 User interface

To perform a simulation, a user specifies a model in a set of Matlab structures. VERTEX automatically initialises the simulation environment, generates the network, and runs the simulation based on these parameters. We divide parameters into five Matlab structures, each holding information about a different aspect of the model and simulation:

- Neuron group properties (for each group: the neurons' compartmental structures, dimensions and positions, electrotonic properties, spiking model parameters, afferent synapse properties)
- Connectivity (for each presynaptic group: number of synapses per layer per postsynaptic group, allowed postsynaptic compartments to connect to contact, axonal conduction speeds, neurotransmitter release times)
- Tissue properties (dimensions, layer boundaries, neuron density, tissue conductivity)
- Recording settings (IDs of neurons to record intracellularly, extracellular electrode positions, sampling rate)
- Simulation settings (simulation length, time-step, number of parallel processes)

A model is initialised by positioning the specified number of neurons from each group within the slice and layer boundaries, pre-calculating distances from the neuron compartments to the virtual electrodes, generating each neuron's connections based on its position, axonal arbour extent in each layer, and expected number of connections, and initialising the synapses. The initialised model can be saved to disk as MAT files, so that it can be reloaded for future simulations, for visualisation, or for analysing the connectivity (see Figure 3.1). We describe these processes in more detail below.

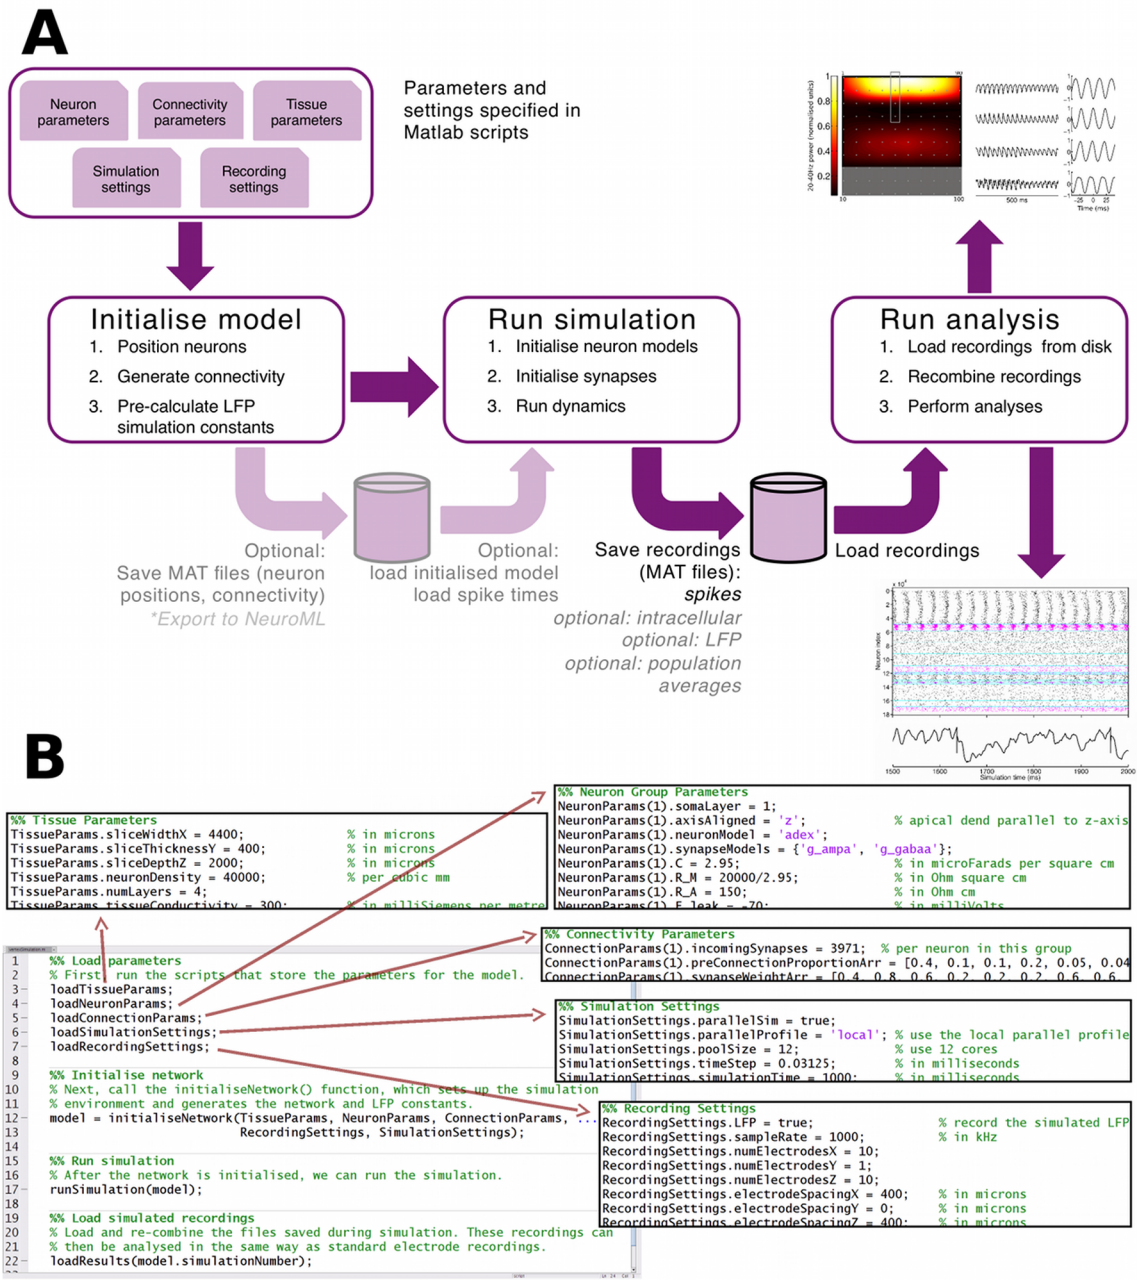

Figure 3.1: Overview of the VERTEX simulation software. (A) Simulation workflow. The user provides parameters as Matlab structures to setup the neuron populations, position them in layers, connect them together, and simulate their dynamics and the resultant LFPs. Functionality to export to NeuroML is currently under development. (B) Example program structure. The main simulation program only requires calls to the initialiseNetwork() function and the runSimulation() function, with the information required to setup the simulation specified in separate script files.

3.3.2 Network initialisation

First, the model is initialised by distributing neurons across parallel processes (if running in parallel), positioning the neurons, setting up the connectivity matrix, and calculating axonal delays. Next, the electrode locations are specified and distances between each electrode and each compartment are calculated, using either the point distance for somas, or the line source distance for dendrites. Pre-calculating the constant values used in the field potential calculations minimises the impact of calculating the LFP during the simulation.

Neurons are positioned by randomly generating a location for their soma compartment within the model and relevant layer boundaries. Each neuron's connections are then generated. This can be done from a presynaptic or postsynaptic perspective. By default, the presynaptic perspective is taken: for each presynaptic neuron, all neurons within its axonal arbour range are considered to be potential targets, and connections are generated one postsynaptic group at a time. In simulations with spatial boundaries, such as the simulated slice we describe below, this enables the calculation of the proportion of each presynaptic neuron's missing axonal arbour according to the model boundaries (see below), and so simplifies the calculation of how many connections each neuron is missing because of slice cutting. Neurons of the same type in models with spatial boundaries will thus have varying out-degree as well as varying in-degree. When boundaries are not specified to reduce the number of connections, the presynaptic perspective means that the out-degree of each neuron in a neuron group is the same (unless specified to have some variance), but the randomness of target selection in the connection algorithm means that the in-degrees will still vary. We therefore also included a connection algorithm that works in the same way as the default method, but takes a postsynaptic perspective, so that neuron groups with constant in-degree can also be generated. This flexibility was important, as in- and out-degree have been shown to have different effects on network dynamics under various conditions (Roxin, 2011).

For models with spatial limits that represent physical boundaries in experiments, such as *in vitro* slice preparations, VERTEX provides a method to calculate each neuron's missing number of synapses due to these boundaries. When a presynaptic perspective is taken, this number can be found for each neuron by calculating the proportion of the axonal arbour remaining within the slice boundaries, and mul-

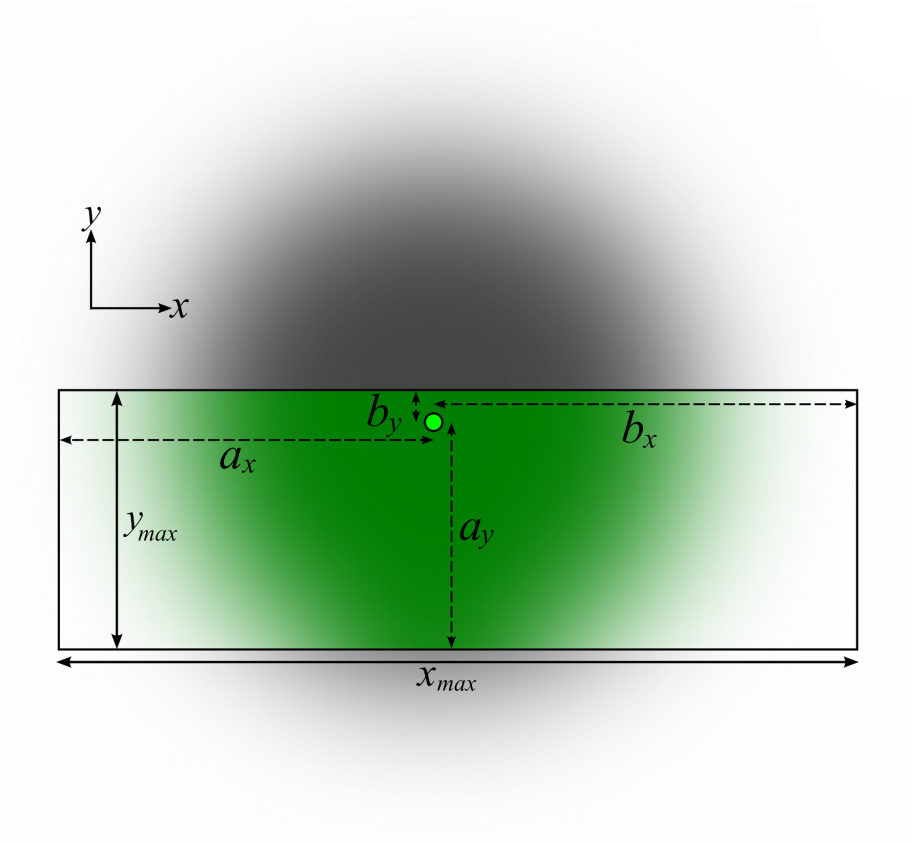

Figure 3.2: Illustration of the effect of slice cutting on a presynaptic neuron's (light green dot) axonal arbour (shaded area). Figure orientation is as if looking down onto the surface of the brain, with slice boundaries indicated by the black bounding box. Connections within the green shaded area remain, but those in the grey shaded areas are removed by slicing.

tiplying this proportion by the expected number of neurons that neuron would have made in an intact brain. The default assumption is that a neuron makes connections within its axonal arbour with a probability that decays away from its soma as a 2D Gaussian in the $x-y$ plane, and the probability in the z-direction is constant within the layer, though the method can be adapted for arbitrary probability distributions. In the case of the 2D Gaussian, the proportion ζ_{li} of neuron i's connections remaining in layer l after slicing is given by

$$\begin{aligned}\zeta_{li} &= \int_{a_{yi}}^{b_{yi}}\int_{a_{xi}}^{b_{xi}} \frac{1}{2\pi\sigma_{li}^2}\exp\left[-\left(\frac{x^2+y^2}{2\sigma_{li}^2}\right)\right]dxdy \qquad (3.1)\\ &= \frac{1}{4}\left\{\left[\mathrm{erf}\left(\frac{a_{xi}}{\sqrt{2}\sigma_{li}}\right)-\mathrm{erf}\left(\frac{b_{xi}}{\sqrt{2}\sigma_{li}}\right)\right]\times\left[\mathrm{erf}\left(\frac{a_{yi}}{\sqrt{2}\sigma_{li}}\right)-\mathrm{erf}\left(\frac{b_{yi}}{\sqrt{2}\sigma_{li}}\right)\right]\right\},\end{aligned}$$

where a_{xi} is the distance from neuron i to the left edge of the slice, b_{xi} the distance to the right edge, a_{yi} the distance to the front edge, b_{yi} the distance to the back edge, σ_{li} is half the arbour radius of i in layer l (see Figure 3.2), and erf the Gaussian error function (this solution is valid provided that a_{xi} and a_{yi} are negative, and b_{xi} and b_{yi} are positive). This is illustrated in Figure 3.2.

Pyramidal neuron dendrites can span several layers above their soma layer, so we decided to allow connectivity to be specified per layer per group, as some datasets (e.g. (Binzegger et al., 2004)) include this information. As all neurons within a population are the same size, but have different soma depths, each neuron's compartments will cross the layer boundaries at different points. This means that each neuron will have different proportions of its compartments in each layer. For simplicity, we ignore this variability for the purposes of generating the connectivity: the layer crossing points for each compartment are found assuming that the soma is at the centre of its layer. We then calculate the membrane area of each compartment in each layer. When connectivity is specified per-layer, the connectivity algorithm calculates a presynaptic neuron's targets one layer at a time. The compartment contacted by the presynaptic neuron in the layer under consideration is chosen randomly, with each compartment having a probability of being selected equal to the membrane area of the compartment in the layer divided by the neuron's total membrane area in the layer. Random targets can be chosen with or without replacement,

allowing multiple connections between two neurons, and autapses can be allowed or prevented. These options can be configured in the simulation parameters, but by default, targets are chosen with replacement and autapses are not allowed. We used these default settings in all our simulations, unless stated otherwise. After the connectivity has been initialised and the relevant connection information exchanged between parallel processes, we calculate the constants that will be used when determining the LFP, so as to minimise the impact of the LFP calculation on simulation speed.

3.3.3 Dynamics simulation

VERTEX allows the simulation of single and multi-compartment neurons, though as described in the previous chapter, multi-compartment neurons are required for LFP simulation. We use the compartmental formulation of the cable equation described in (Gerstner and Kistler, 2002) to model the passive dynamics of each neuron. The user specifies the compartmental structure, axial resistivity, membrane resistivity and membrane capacitance when the model is initialised. VERTEX then calculates the inter-compartmental resistance, axial resistance, membrane resistance and capacitance of each compartment from these user-specified parameters. The compartmental coupling resistance between compartments i and j $R_{ij} = (R_{a,i} + R_{a,j})/2 = 1/g_{ij}$ (Gerstner and Kistler, 2002), with the axial resistance of each compartment $R_a = lR_A/\pi r^2$. The membrane potential of each compartment i then evolves according to equation 3.2:

$$C_{m,i}\frac{dv_i}{dt} = -g_{l,i}(v_i - E_l) - \sum_j g_{ij}(v_i - v_j) + I_i, \tag{3.2}$$

where $C_{m,i}$ is the compartment's membrane capacitance, v_i is compartment i's membrane potential, $g_{l,i}$ is its leak conductance (= reciprocal of membrane resistance), E_l is the leak reversal potential, g_{ij} are the conductances between compartment i and its neighbours j, v_j is the membrane potential of the jth connected compartment, and I_i is the input current to the compartment (including synaptic and injected currents).

We limited ourselves to models with passive dendritic compartments, for several reasons: firstly for simplicity, as estimating parameters for active conductances

is notoriously difficult, especially if attempting to fit extracellular as well as intracellular measurements (Gold et al., 2007, 2006); secondly because the reduced compartmental models we were using had not been tested with active conductances, including these conductances on the reduced morphology would have potentially unrealistic effects on the extracellular potential; thirdly, because active conductances add further run-time costs. Models with active conductances are possible to implement in VERTEX, but are not included by default.

For spiking neural network simulations, some mechanism is required by which a neuron "decides" when to fire. We therefore implemented a version of the passive compartmental model that also contained an adaptive exponential (AdEx) (Brette and Gerstner, 2005) firing mechanism at the soma compartment as VERTEX's standard spiking neuron model. The AdEx model can reproduce most of the dynamical features exhibited by cortical neurons (Naud et al., 2008), all its parameters have a direct biological correlate (Gerstner and Brette, 2009) making the model easy to interpret and modify in light of new experimental data, its sub-threshold behaviour is realistic (Badel et al., 2008), and its bifurcation structure is well characterised and is the same as the commonly used Izhikevich model (Naud et al., 2008; Touboul and Brette, 2008). Crucially, it can easily be extended to include passive dendrite compartments (Clopath et al., 2007; Gerstner and Brette, 2009) as it is formulated as a conductance-based model. It is also much faster to simulate than a Hodgkin-Huxley type model (Izhikevich, 2004). The somatic membrane potential v_s in the AdEx compartmental neuron then evolves according to equations 3.3 and 3.4:

$$\begin{aligned} C_s \frac{dv_s}{dt} &= -g_{leak,s}\left(v_s - E_{leak}\right) - \sum_j g_{sj}\left(v_s - v_j\right) + I_s \\ &\quad + g_{leak,s}\Delta_t \exp\left(\frac{v_s - V_t}{\Delta_t}\right) - w, \end{aligned} \tag{3.3}$$

$$\tau_w \frac{dw}{dt} = \alpha\left(v_s - E_{leak}\right) - w, \tag{3.4}$$

where Δ_t is a constant defining the spike steepness, V_t is the instantaneous threshold potential, w is a current representing the combined slow ionic currents, τ_w is the time constant of the slow current w, and α is the scale factor of the slow current (Brette and Gerstner, 2005) (other variables and parameters are as in equation 3.2, for soma compartment s rather than dendritic compartment i) A neuron is said to have fired a

spike when the membrane potential reaches a value v_{cutoff}, which we always set to be equal to $V_t + 5$mV. We found that having v_{cutoff} set too high could cause the membrane potential to explode to infinity even within a single time-step due to the exponential term. Setting v_{cutoff} to $V_t + 5$mV was recommended by Marc-Oliver Gewaltig on the NEST mailing list as allowing the exponential term to influence the upstroke of the spike without causing numerical issues (archive at http://ken.brainworks.uni-freiburg.de/cgi-bin/mailman/private/nest_user/ for date 2012-04-26). When a spike is fired, i.e. $v_s \geq v_{cutoff}$, the following substitutions are made:

$$\begin{aligned} v_s &\rightarrow v_{reset}, \\ w &\rightarrow w + \beta, \end{aligned}$$

where v_{reset} is the membrane potential to which v_s returns after a spike, and β is the instantaneous increase in the value of the slow current w after a spike (Brette and Gerstner, 2005).

The exponential term and slow current create additional membrane currents in the otherwise passive cable models, so have an influence on the extracellular potential. In particular, the instantaneous reset of the soma potential causes high frequency spikes in the signal that are not physiologically realistic, and the slow current w also contributes to the lower frequencies with unknown physiological relevance. Depending on the number of neurons, their firing rates, the synaptic weights, and the level of synchrony in the network, these contributions can often be negligible compared to the synaptic contribution to the extracellular potential (see Chapter 4, Figure 4.8). However, in cases where these currents make a significant contribution to the simulated LFP, the membrane currents can be recalculated to include only the synaptic contributions, by re-running the simulation with purely passive neurons and using the previously generated spike times to drive the synaptic activity (see "Spike loading" below).

The input current I_i to compartment i represents the sum of synaptic currents $I_{syn,i}$ and any other externally applied currents $I_{ext,i}$. We provide both current-based and conductance-based synapse models, with exponential decays, and used both types in our simulations. For current-based synapses, incoming spikes affect the synaptic current directly causing it to increase instantaneously by the synaptic

weight ($I_{syn,i} \rightarrow I_{syn,i} + W$), after which it decays exponentially:

$$\frac{dI_{syn,i}}{dt} = -\frac{I_{syn,i}}{\tau_{syn}}, \tag{3.5}$$

where τ_{syn} is the synaptic decay constant. For conductance-based synapses, the compartment's synaptic conductance $g_{syn,i}$ is increased by the synapse weight ($g_{syn,i} \rightarrow g_{syn,i} + W$) rather than the current being modified directly, and the synaptic current is calculated as

$$I_{syn,i} = -g_{syn,i}(v_i - E_{syn}), \tag{3.6}$$

where E_{syn} is the reversal potential for that type of synapse. Again, the conductance decays exponentially:

$$\frac{dg_{syn,i}}{dt} = -\frac{g_{syn,i}}{\tau_{syn}}. \tag{3.7}$$

Contributions from each synapse onto a compartment sum linearly, so only one variable for the synaptic current (and one for the conductance, if conductance-based synapses are used) is required, reducing calculation time and memory usage.

We provide several types of currents and conductances that can be applied to individual compartments or whole neurons, including constant current injections and randomly fluctuating currents or conductances. We most commonly use an Ornstein-Uhlenbeck process to model random background input, as used to model *in vivo* conductance (Destexhe et al., 2001) or current (Arsiero et al., 2007) fluctuations. The update rule for the Ornstein-Uhlenbeck process is the exact discretisation method described in (Gillespie, 1996). In the case of a fluctuating conductance $g_{ext,i}$, the update rule is:

$$\begin{aligned} g_{ext,i}(t+\delta t) &= g_{ext,i}(t) + \left(1 - \exp\left(\frac{-\delta t}{\tau_{ext}}\right)\right) \times (m_{ext} - g_{ext,i}(t)) \\ &\quad + \sqrt{1 - \exp\left(-2\frac{\delta t}{\tau_{ext}}\right)} \times S_{ext} \times N_{ext}(0,1), \end{aligned} \tag{3.8}$$

where δt is the length of the time-step, τ_{ext} is the noise correlation time constant, m_{ext} is the mean current value, S_{ext} is the standard deviation and $N(0,1)$ is a normally distributed random number. The current is then calculated as

$$I_{ext,i} = g_{ext,i}(v_i - E_{ext}), \tag{3.9}$$

where E_{ext} is the externally applied conductance's reversal potential. In the case of a fluctuating current, equation 3.9 is used to update the value of the current directly (the $g_{ext,i}$ terms are simply replaced by $I_{ext,i}$).

LFPs can be calculated online (during the simulation) or offline (after the simulation has finished) by storing all compartments' membrane currents. Online calculation slows the simulation and loses information regarding the LFP contribution of each individual compartment, but greatly reduces the amount of disk space required as only one time series per electrode, rather than one time series per compartment, is saved. Calculating LFPs offline also increases the time required to load the saved data from disk. LFPs are found by summing the membrane currents of each compartment weighted by distance from the electrode tips, using the line-source method for dendritic compartments and point-source method for somatic compartments as described in the previous chapter. We use the method described in (Holt, 1998) for calculating the potential due to a line current source, which reduces round-off error when $h > 0$ and $l > 0$:

$$\Phi = \frac{I}{4\pi\sigma\Delta s} \begin{cases} \log \frac{\sqrt{h^2+\rho^2}-h}{\sqrt{l^2+\rho^2}-l} & \text{for } h < 0,\ l < 0 \\ \log \frac{\left(\sqrt{h^2+\rho^2}-h\right)\left(\sqrt{l^2+\rho^2}+l\right)}{\rho^2} & \text{for } h < 0,\ l \geq 0 \\ \log \frac{\sqrt{l^2+\rho^2}+l}{\sqrt{h^2+\rho^2}+h} & \text{for } h \geq 0,\ l \geq 0, \end{cases} \tag{3.10}$$

where the numerator and denominator in equation 2.9 in Chapter 2 have been multiplied by $\sqrt{l^2+\rho^2}+l$ when $l \geq 0$, or $\sqrt{h^2+\rho^2}+h$ and $\sqrt{l^2+\rho^2}+l$ when $h \geq 0$ & $l \geq 0$ (Holt, 1998). This method is also implemented by LFPy. Also similar to LFPy, we set a minimum distance that a current source can be to an electrode, so that the electrode cannot be "inside" a compartment. This minimum distance can be specified as a parameter set by the user. The membrane current $I_{mem,i}$ of compartment i is just the negative of the axial current $I_{ax,i}$ entering the compartment (Johnston and Wu, 1995):

$$I_{mem,i} = -I_{ax,i} = \sum_j g_{ij}\left(v_i - v_j\right). \tag{3.11}$$

3.3.4 Numerical methods

We used a second order Runge-Kutta method to solve all the differential equations. The forward Euler method is the simplest and fastest numerical integration method, and is equivalent to a first-order Runge-Kutta method. The global error for the Euler method is proportional to the length of the time-step. We chose the second order Runge-Kutta method (specifically, the midpoint method) because it remains relatively easy to implement compared to the Euler method, while reducing the global error to be proportional to the time-step squared (Press et al., 2007). It is more computationally intensive, but not prohibitively so. It is also a commonly used method in neuronal network simulations (Henker et al., 2011), and is implemented as a standard integration method in the Brian simulator (Goodman and Brette, 2009). The size of the time-step can be specified by the user; in all our simulations, we used a time-step of 0.0325ms. Smaller time-steps did not qualitatively affect the results for the simulations we ran, while larger time-steps resulted in numerical instability when using the compartmental AdEx neuron models. For an investigation into the accuracy of various numerical methods commonly used in neural simulators, see (Henker et al., 2011).

3.3.5 Vectorisation methods

Interpreted languages like Matlab and Python incur a performance overhead with every line of code executed, as the interpreter must first translate the line to machine code before executing it. For intensive computations this performance overhead can be extremely costly, especially in large systems. For example, to update all the membrane potential variables in a time-step, the programmer would traditionally write a loop to update each variable in the simulation one at a time, thus incurring the interpretation overhead at each step of the loop. Even for relatively small simulations, this performance overhead can become prohibitively large.

Matlab (and the NumPy extension for Python) has a built-in programming construct to help alleviate this performance issue. Instead of individually applying a particular function to every element of a matrix, one can apply the function to the entire matrix in a single step. As the built-in mathematical functions are highly optimised, they run at approximately the same speed as equivalent functions in

compiled C code. By reducing the interpretation overhead to a single command that runs the optimised function over the whole matrix instead of incurring the overhead for every matrix element in a loop, execution time can be dramatically reduced. The performance improvements from vectorisation improve with array size, too: the interpretation overhead is roughly constant, so the proportional time for interpretation decreases with increasing array size (Brette and Goodman, 2011).

We made use of the vectorised algorithms and data structures outlined in (Brette and Goodman, 2011). This paper describes vectorised solutions for common tasks in neuronal network simulation including updating the neuronal state matrices, initialising network connectivity, and propagating spikes including transmission delays. However, updating the state matrices for compartmental neuron models is not covered. We therefore developed a vectorised method to handle this. We do not claim that the method is optimal, but it is certainly faster than iterating through compartments individually.

The method assumes that the compartmental structure of a neuron forms a binary tree, so each compartment has a maximum of three neighbours: its parent and up to two children. The somatic compartment is the root of the tree, so has no parent compartment. The tree is specified as an array listing the parent compartment of each other compartment. During initialisation, VERTEX converts this array into three matrices: the first containing the mapping of compartments to their parents, the second the mapping of compartments to their first children, and the third the mapping of compartments to their second children (see Figure 3.3 c). The conductances between the compartments are also calculated and stored in three arrays corresponding to the parent, child 1 and child 2 matrices. The current I_{ij} at compartment i due to compartment j is then $-g_{ij}\,(v_i - v_j)$. By indexing the membrane potential matrix using the parent, child 1 and child 2 matrices, we can calculate all the current flows between the compartments in three vectorised steps with the values for i and j stored on the parent, child 1 and child 2 matrices. This is illustrated in figure Figure 3.3.

3.3.6 Parallel programming methods

Vectorisation can dramatically improve performance, but large-scale simulations can still take a long time through the sheer number of calculations that must be executed.

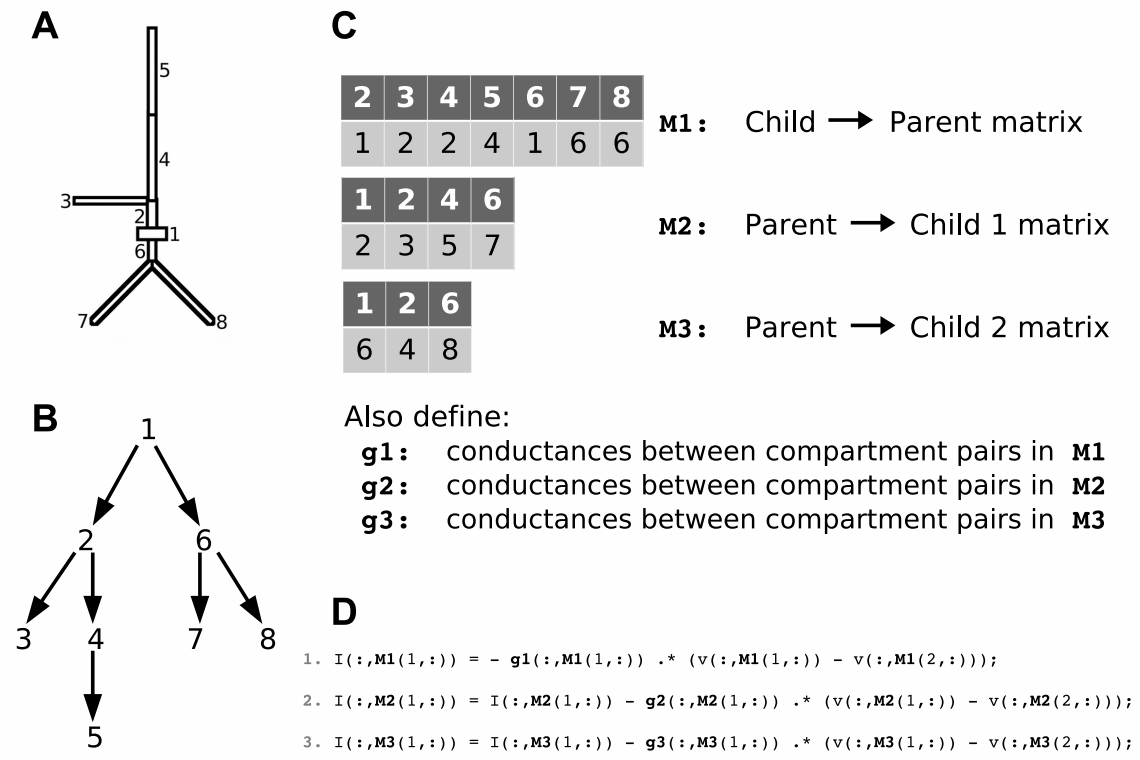

Figure 3.3: Illustration of vectorisation method for updating currents between compartments. (A) Example neuron structure with numbered compartments (in this case, a layer 2/3 pyramidal neuron with 8 compartments). (B) Binary tree representation of the compartment structure. (C) The three matrices (`M1`, `M2`, `M3`) listing the connections between compartments. During model initialisation, three matrices (g1, g2, g3) containing the conductances between the compartment pairs defined in `M1`, `M2` and `M3` are also calculated. (D) Matlab code for updating the inter-compartment currents. `I`, `v`, `g1`, `g2` and `g3` are matrices with each row representing a neuron and each column representing a compartment in the neuron. `I` stores the total current at each compartment due to its neighbouring compartments, and v stores the membrane potential of each compartment.

Parallel programming can further reduce simulation time by splitting these calculations across many processors so that each one performs a subset of the calculations simultaneously. Matlab's Parallel Programming Toolbox provides an easy way to define parallel simulations. We used the Single Program Multiple Data (SPMD) formulation to distribute the variables for each neuron group among the parallel processes (or "labs" as they are referred to in Matlab). SPMD, as its name suggests, allows the same code to execute on the different variables spread across the labs. Crucially, it also allows labs to communicate data with each other. This formulation is ideal for simulating spiking networks: the main simulation loop runs the same in each lab, but operates on a different subset of neurons. The ability to send and receive data between labs allows neurons to communicate with each other, whichever lab they happen to be on.

We used this formulation to implement the parallel algorithms described in (Morrison et al., 2005) for neuronal communication. These methods assume that spikes are discrete events that are locked to a time grid, which we set to be the same as the integration time-step. They minimise the communication overhead by storing synapse information (delays, postsynaptic neuron IDs and compartment IDs) on the postsynaptic side, so only spiking presynaptic IDs and timestamps need to be exchanged between processes.

The spike exchange algorithm works as follows (Morrison et al., 2005). At the start of a simulation, a presynaptic spike buffer is created. This buffer consists of a circular array with a length equal to the number of time-steps of the smallest axonal delay time in the model, and a variable that points to the current index to be accessed in the circular array. Each time-step, the pointer variable is incremented after the IDs of the neurons that have reached spiking threshold during that time-step are stored in the currently indexed array element. Once the last element of the array is reached, the buffer is sent to all other parallel processes, and the indexing variable completes the circle back to the first array element. After the spikes have been exchanged, they are then put into postsynaptic buffers, which assign each presynaptic spike to the correct postsynaptic compartment, to be delivered after the relevant axonal delay time. Because we know when the presynaptic neuron spiked given the length of the presynaptic spike buffer, we can subtract the relevant time since spiking from the delay time on the postsynaptic side. This method minimises

both the amount of information to exchange (just presynaptic neuron IDs) and the number of exchanges (once per minimum axonal delay time, rather than every time-step).

We also use the same algorithm as in (Morrison et al., 2005) to arrange the sequence in which processes do the information exchange to ensure minimal wait times. This is the Generalised Pairwise Exchange Algorithm (GPEX), listed below (Morrison et al., 2005; Tam and Wang, 2000). Each lab must send its spike list to every other lab, and receive back every other lab's list of spikes. The order in which the labs communicate is decided by the GPEX algorithm, which ensures that all pairs of labs communicate in a minimum number of steps, and without bottlenecks. This algorithm runs independently on each lab.

3.3.6.1 Simulation speed

To give the user an idea of the performance compared to LFPy, we performed equivalent simulations using layer 5 Bush pyramidal neurons in LFPy and in VERTEX (no synapses, one random fluctuating current per neuron, 0.0325ms step size & sample rate). LFPy took ∼260 minutes to simulate the LFP from 10 000 neurons at 50 electrode points, while VERTEX running in serial mode took ∼14 minutes to simulate the LFP from 10 000 neurons at 50 electrode points (both running on a 2.66 GHz Intel Core i7-920 desktop computer with 6 GB RAM). While this performance improvement is important for our purposes, it should be noted that LFPy is designed to simulate extracellular potentials from single cells rather than large populations. As it gives the user all the power and flexibility of the NEURON simulator (including access to a large selection of single neuron models from ModelDB - http://senselab.med.yale.edu/modeldb/ - or NeuroMorpho.org (Ascoli et al., 2007)) to achieve this, it remains the superior tool for smaller-scale studies where more fine-grained control is required.

To show how VERTEX's performance improves in parallel mode, we compared the run times for two network models, one large (123 517 neurons with on average 1 835 synapses per neuron) and one small (9 881 neurons with on average 256 synapses per neuron), using VERTEX on a single multicore server. The maximum number of labs available on a single computer is currently limited to 12 by Matlab's licensing restrictions, but further cores could be recruited by using the Matlab Distributed

Algorithm 3.1 Generalised Pairwise Exchange Algorithm, adapted from (Tam and Wang, 2000). `p` is the number of labs, `myid` is the ID of the local lab, and `k` is the number of data packages to be exchanged. The Matlab Parallel Programming Toolbox provides the labSendReceive() function, which we use to implement lines 7-8 (this function simultaneously sends and receives the specified data between the two specified labs).

```
1  round = odd(p) ? p : p-1
2  for i = 1 to round do
3     partner = EdgeColor(i,myid,p)
4     if (partner = -1)
5        I am idle in this step
6     else
7        send_data(data, partner)
8        received[partner] = receive_data(partner)
9     endif
10 endfor
11
12 EdgeColor(i, myid, p)
13    S = odd(p) ? p : p-1
14    if (myid < S')
15       v = (i + S - myid) mod S
16    else
17       v = odd(i) ? (((i+S)/2) mod S) : i/2
18    endif
19    if (odd(p) AND v = myid)
20       partner = -1
21    else if (v = myid)
22       partner = S
23    else
24       partner = v
25    endif
26    return partner
```

Computing Server, either in a virtual cluster on the single multicore machine or by using a physical cluster. Each model had two populations: layer 5 pyramidal (P5) neurons and layer 5 basket (B5) interneurons. Spike rates in each small model (large model) simulation were ∼6Hz (∼7Hz) and ∼24Hz (∼31Hz) for the P5 and B5 neurons, respectively. The large model shows linear speed-up with increasing number of cores for model initialisation and close-to-linear speed-up in simulation time (Figure 3.4). The speed-up for the small model is sub-linear: as the interpretation overhead for a vectorised operation on a small matrix is the same as on a large matrix, this overhead starts to dominate the calculation times below a certain number of neurons (Brette and Goodman, 2011). Therefore, splitting already small neuron state matrices between more processes does not significantly improve performance. This limit is not reached in larger models.

Figure 3.4 also shows how increasing the number of virtual electrodes affects simulation speed. Using more electrodes affects initialisation times proportionally more than run times, in both the large and the small model. The proportional impact on initialisation time from adding electrodes was greater in the small model than in the large model. This is because the large model not only has more neurons, but also more synapses per neuron. The increase in time spent connecting the neurons is proportional to the number of synapses, while the increase in time spent calculating constants for the LFP measurements is proportional to the number of compartments (roughly proportional to the number of neurons).

3.3.7 Data recording and analysis

When the simulation is run, recordings (membrane potentials, LFPs, spike times, synaptic currents) are automatically saved to disk at user-specified time intervals (we tended to save every 100ms - 200ms, so that memory usage did not grow too large). The simulation run can be performed in serial or parallel (requires Matlab Parallel Computing Toolbox). After the simulation is finished, these files are loaded and recombined for analysis.

Network dynamics can be simulated directly by providing the model neurons with a spiking mechanism. Alternatively, previously generated spike times (arbitrarily specified, or generated by spiking neurons in VERTEX, or by another simulator) can be imported into the simulation. Neurons are then specified with purely passive

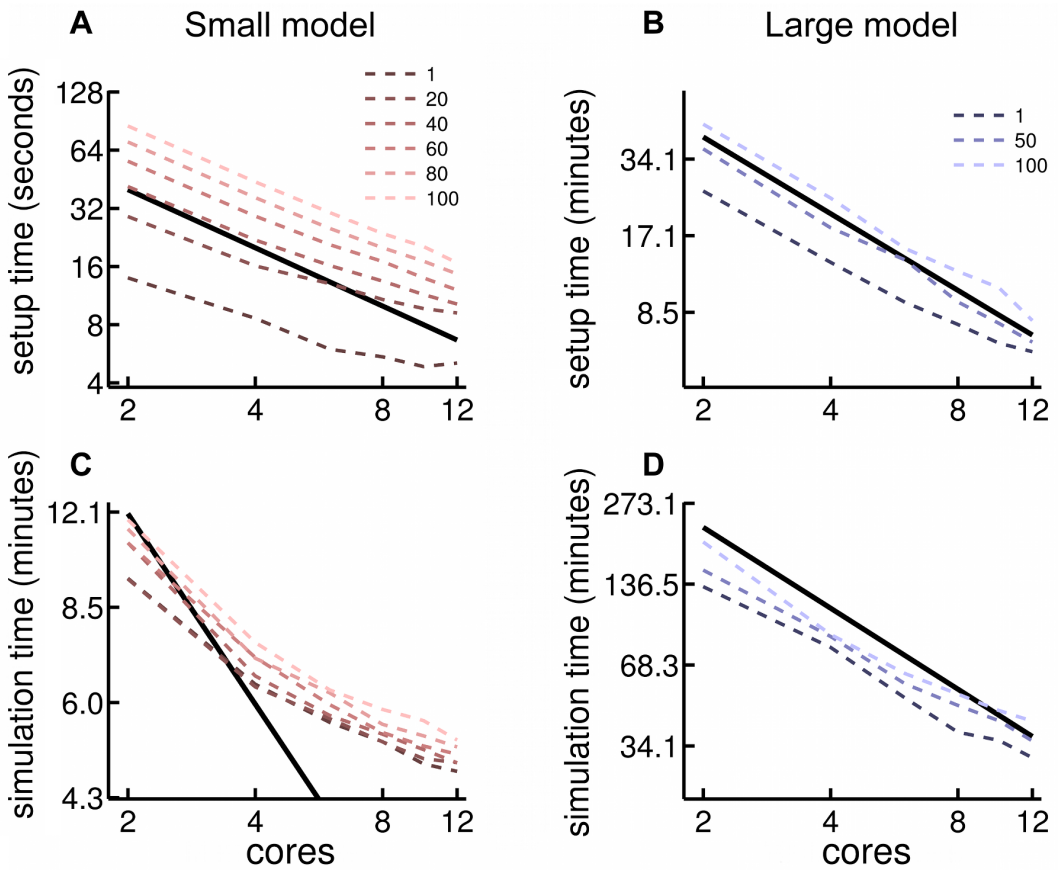

Figure 3.4: Parallel simulation performance with increasing numbers of Matlab workers (i.e. parallel processes). Top row: model initialisation times for (A) the 9 881 neuron model and (B) the 123 517 neuron model. Bottom: simulation times for 1 second of biological time for (C) the 9 881 neuron model and (D) the 123 517 neuron model. Thick black lines indicate linear speed scaling; legends indicate the number of electrodes used in each simulation run. The sub-linear speed-up in the small model is due to the decreasing relative performance influence of code vectorisation for smaller matrices (see Results).

membrane dynamics. This is similar to the approach used in (Lindén et al., 2011) to link spiking output from a cortical model implemented in the NEST simulator (Gewaltig and Diesmann, 2007) to their LFP generating model implemented in LFPy. However, VERTEX simplifies the simulation of the LFP due to a specific network structure, as imported spikes are delivered to target neurons according to the generated connectivity matrix rather than pre-assigned to postsynaptic targets. By contrast, in (Lindén et al., 2011) the spikes from NEST-simulated neurons were considered as external input to the neurons in the LFPy simulation, so were delivered to synapses without a connectivity model within the LFPy-simulated population. VERTEX simplifies modelling the LFP resulting from intrinsic network dynamics, when connectivity is known or when different spatial connectivity models are to be tested. We used the spike import feature to run the control experiment to confirm that the AdEx spiking mechanism has a negligible impact on the simulated LFP (see Chapter 4).

3.4 Summary

In this chapter, we described the VERTEX simulation tool, which we developed to address our need for a simulator that would allow us to investigate LFP generation in spiking neuronal network models. VERTEX includes functionality for generating spatially constrained networks of several neuron populations, whose parameters are easily specified in Matlab structures. "Virtual electrodes" can be positioned at arbitrary locations in the model to simulate the LFP generated by the network. Parallel computing and code vectorisation allows VERTEX to simulate network activity and LFPs in reasonable time.

In the next chapter, we describe the implementation and simulation results from a neocortical network model. The model illustrates the ability of VERTEX to scale to large model sizes. We also show how simulations using VERTEX can be used in conjunction with experimental extracellular recordings.

Chapter 4

Simulating the LFP during persistent gamma oscillations

4.1 Introduction

In the previous chapter, we described the VERTEX simulation tool. We developed VERTEX to allow us to construct and simulate dynamics in large networks of compartmental neuron models, and to simulate the local field potentials arising from this activity. In this chapter, we describe our implementation of a neuronal network model in VERTEX, representing a preparation of neocortical tissue *in vitro*. We developed the model to gain insights into the experiments being performed *in vitro* by our collaborators at the Newcastle University Institute of Neuroscience. These experiments examined the extracellular potentials in macaque neocortical slices exhibiting gamma frequency oscillatory activity induced by bath application of the glutamate receptor agonist kainic acid (kainate). They used Utah multi-electrode arrays, which cover a 3.6 mm by 3.6 mm area with 100 electrodes, spaced in a square grid. We designed the model with these dimensions in mind, so that we could simulate extracellular recordings using a virtual Utah array, allowing a direct comparison to be made between the experimental and simulated results.

The *in vitro* persistent gamma model also had some benefits for simulating the LFP. First, the theory of how neocortical persistent gamma arises *in vitro*, and how individual neurons participate in the network oscillation, has been comprehensively documented (Ainsworth et al., 2011; Bartos et al., 2007; Buhl et al., 1998; Cunningham et al., 2003, 2004b; Draguhn et al., 1998; Fisahn et al., 1998; Roopun et al.,

2008; Traub et al., 2005a,b; Whittington et al., 2011, 1995). We could therefore fit the spiking activity of the network to the experimentally observed spiking patterns and simulate the resulting LFPs across the network. Second, the slice preparation ensures that all synapses are local. The MEA recordings are therefore influenced only by the local circuit dynamics, so we did not need to worry about including synaptic input from other brain regions in the model: the slice edges provide natural spatial boundaries for what needs to be included in the model. Third, persistent gamma is known to depend on rapid synaptic interactions, rather than intrinsic oscillatory currents in individual neurons. We could therefore assume that the majority of the experimental signal is synaptically generated, so our LFP model incorporating only synaptic and passive return currents would be adequate for simulating the LFP in this activity regime.

In the first part of this chapter, we describe the experimental persistent gamma oscillation model: how it is generated, the neuronal spiking activity that creates it, and experimentally observed features of the LFP. In the second part, we describe our neocortical slice model and compare our modelling results with experimental data.

4.2 Persistent gamma frequency oscillations *in vitro*

4.2.1 Why model persistent gamma oscillations?

In Chapter 1, we discussed the presence of oscillations in EEG and LFP recordings. Particular oscillations correlate with particular behavioural states, or show alterations in various neurological disorders. Much of the work done to understand the mechanisms of these oscillations has been performed in brain slice preparations *in vitro* (Traub and Whittington, 2010). The *in vitro* persistent gamma oscillation model first described in (Fisahn et al., 1998) has been particularly useful in unravelling the mechanisms involved in generating gamma frequency rhythms in the brain: it provides a stable model of an oscillation that can be studied over a period of several hours, and the firing patterns of the different neuron groups appear to correspond well to firing during gamma oscillations *in vivo* (Csicsvari et al., 2003; Hájos and Paulsen, 2009). Understanding the generation of the LFP during this well

characterised *in vitro* oscillation could therefore shed light on the neuronal activity during gamma oscillations *in vivo*.

4.2.2 Persistent gamma oscillation mechanism

Gamma frequency oscillations can be induced in cortical slices *in vitro* by bath application of the glutamate receptor agonist kainate (Cunningham et al., 2003, 2004b,a; Traub et al., 2005a), or cholinergic agonist carbachol (Fisahn et al., 1998; Buhl et al., 1998). These kinds of oscillation are frequently referred to as persistent gamma oscillations as they remain stable over several hours (Whittington et al., 2011), distinguishing them from shorter periods of gamma activity evoked by tetanic stimulation or localised glutamate application. The cellular and synaptic mechanisms underlying persistent gamma are well characterised and have been studied in detail in hippocampus (Atallah and Scanziani, 2009; Buhl et al., 2003; Draguhn et al., 1998; Ferguson et al., 2013; Oren et al., 2006; Traub et al., 1996a; van der Linden et al., 1999; Whittington et al., 2001, 1997b,a; Traub et al., 2003b; Whittington et al., 1995), entorhinal cortex (Cunningham et al., 2003, 2004a; Dickson et al., 2000; van der Linden et al., 1999) and neocortex (Ainsworth et al., 2011; Cunningham et al., 2004b; Haenschel et al., 2000; Traub et al., 2005a; Whittington et al., 1997b, 1995). The proposed mechanism generating the oscillation involves the excitation of the neuronal network through spiking in pyramidal cell axons. Each pyramidal cell axon is coupled to other pyramidal axons by gap junctions, which allow spikes to cross between the axons if the junction resistance is low enough. The active channels in each axon allow each spike to propagate through many other axons in the gap-junctionally connected axonal plexus, so that a single spike will cause glutamate to be released at many pyramidal axon terminals. This provides continuous excitation to the cells in the network, as pyramidal cells form synapses with most other cell types (Binzegger et al., 2004).

The continuous excitation provided by the active axonal plexus drives a gamma oscillation by means of synchronous firing of inhibitory interneurons; in particular, fast-spiking basket cells. The exact mechanism by which the interneurons generate the rhythm, and the dependence on different excitatory cell populations, varies between brain regions. For example, a study in rat neocortical slices suggested that persistent gamma in this region was dependent upon the existence of a small pop-

ulation of fast rhythmic bursting (chattering) pyramidal cells in superficial layers, which inject bursts of action potentials into the axonal plexus at gamma frequency (Cunningham et al., 2004b). In hippocampus this cell type is not required (Cunningham et al., 2004b); a possible cause of the continued the axonal plexus activity in hippocampus is the spontaneous release of gamma-Aminobutyric acid (GABA) at interneuron axon terminals, which can have an excitatory effect on pyramidal cell axons (Traub et al., 2003b). Additionally in auditory neocortex, different gamma generating networks exist in different layers, and are activated differentially by different concentrations of kainate (Ainsworth et al., 2011).

The key aspect of persistent gamma in all the different preparations is fast inhibition (Bartos et al., 2007; Fisahn et al., 1998; Mann et al., 2005). In each model, fast spiking interneurons fire synchronously, in phase with the oscillation as measured in the LFP Hájos et al. (2004). These interneurons make many connections to the perisomatic regions of pyramidal cells, providing them with strong phasic inhibition (Kawaguchi and Kubota, 1997). They are also recurrently connected, so receive inhibition immediately after their synchronous firing. The oscillation frequency is then dependent upon (amongst other factors such as strength of external drive and membrane time constants (Economo and White, 2012)) the decay time constant of the inhibitory synapses (Fisahn et al., 1998), as this defines the amount of time that the strong inhibitory current lasts and so the time window during which a cell is prevented from spiking. If we consider this crucial dependence on fast phasic inhibition to be the defining factor of persistent gamma, then we can create a model of the oscillation that captures this aspect of the oscillatory mechanism while ignoring many of the variable details, including the unknown parameters of the axonal plexus and the dependence on chattering cells in neocortex. Simplified mathematical models have shown that persistent gamma-like oscillations can be generated in networks of excitatory and inhibitory neurons even when the details of the excitatory drive from activity in the axonal plexus are ignored, and the drive is approximated by independent, random synaptic inputs or randomly fluctuating currents (Ainsworth et al., 2011; Börgers and Kopell, 2005; Economo and White, 2012).

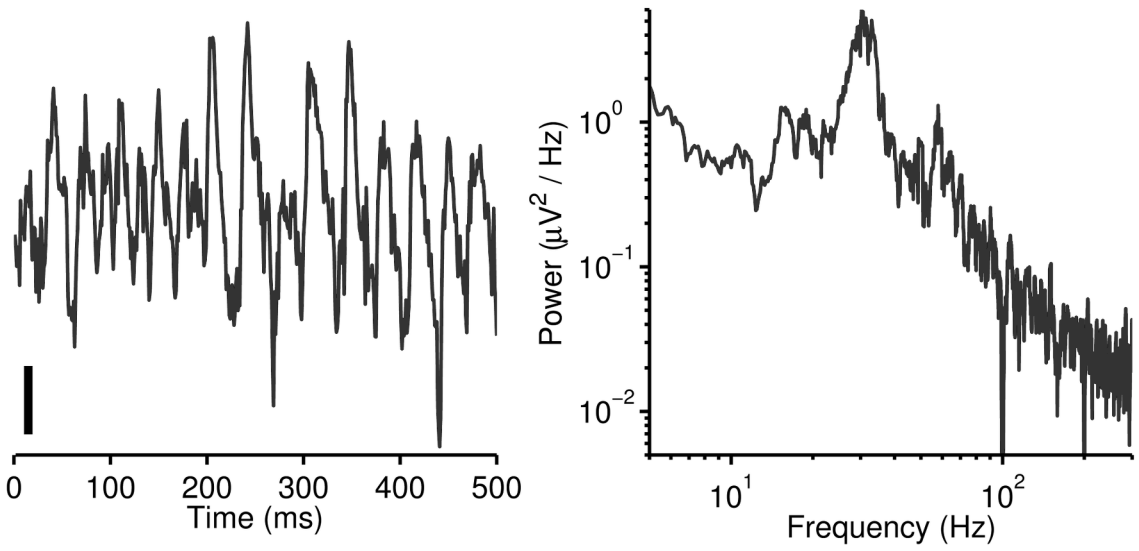

Figure 4.1: Example of an extracellular electrode recording (left) performed in macaque auditory neocortex *in vitro*, showing a strong gamma frequency oscillation as revealed by the large peak at 30 Hz in the power spectrum plot (right). The electrode was positioned in cortical layer 3. The gamma oscillation was induced by bath application of the glutamate agonist kainate. Scale bar: 10 microVolts. Raw data provided by Matthew Ainsworth at the Newcastle University Institute of Neuroscience.

4.2.3 LFP features during persistent gamma

The most obvious feature of persistent gamma is the very clear rhythmic component in the signal, often at a frequency of between 30Hz and 40Hz. This signal is stable for many hours, but varies across space. Several studies have investigated the spatial profile of gamma oscillations, usually looking at the variation with cortical depth using laminar electrode arrays (Dickson et al., 2000; van der Linden et al., 1999). Recent recordings made by our collaborators using Utah MEAs also allow the investigation of lateral variations in the LFP (Ainsworth, 2013).

One spatial feature of neocortical gamma oscillations is of particular note. A phase inversion in the LFP is observed when the electrode is moved down from the cortical surface through the layers towards the white matter boundary. This inversion usually occurs somewhere between layers 1 and 3, and has been observed in guinea pig whole-brain preparations (Dickson et al., 2000; van der Linden et al., 1999), in rat entorhinal cortex *in vitro* (Cunningham et al., 2003) and in recordings made in freely behaving rats (Csicsvari et al., 2003). The reason for this phase inversion was investigated by Cunningham et al. (2003), who observed a phase change between layers 2 and 3. They measured inhibitory postsynaptic potentials (IPSPs)

in layer 2 stellate cells and layer 3 pyramidal cells simultaneously with the LFP, showing that IPSPs in both layer 3 cells and layer 2 cells were in anti-phase with the LFP, while IPSPs in layer 3 cells were in phase with the layer 2 LFP. This suggested that neurons in layers 2 and 3 shared the same source of inhibitory input, and that the phase inversion was caused by a separation of current sinks and sources rather than separate gamma generating networks (localised perisomatic inhibitory currents result in an opposite return current in the apical dendrites of the pyramidal neurons).

4.2.4 Individual neuron spiking during persistent gamma

The spiking behaviours of different types of neurons during persistent gamma oscillations have been confirmed in many previous experiments. Fast spiking interneurons fire synchronously on most or all of the gamma cycles, while pyramidal cells fire more sparsely and tend to be less strongly locked to the oscillation phase (Hájos et al., 2004). Other interneuron types also fire in phase with the oscillation, but at lower rates than the fast-spiking interneurons. The low firing rate of pyramidal neurons appears to be a good representation of the *in vivo* case, in which pyramidal neurons spike at much lower frequency than the network oscillation (Csicsvari et al., 2003; Hájos and Paulsen, 2009).

Using this information on spike timing, and combining it with a connectivity model so that the postsynaptic effects of a spike make a suitable contribution, allows us to simulate the LFP resulting from this spiking activity: given the firing pattern of each neuron, the connectivity between neurons, and their positions in space, what is the theoretically predicted LFP at a particular location? In the next section, we describe the network model we implemented to integrate the relevant information for studying the spatial LFP properties during persistent gamma oscillations.

4.3 Neocortical Slice Model

The novel use of the Utah MEA allowed our experimental collaborators to record LFPs at many locations simultaneously. To model the recordings from these experiments, we not only needed to consider the dynamics of the individual neurons and synapses, but also the connectivity between neurons and the spatial organisa-

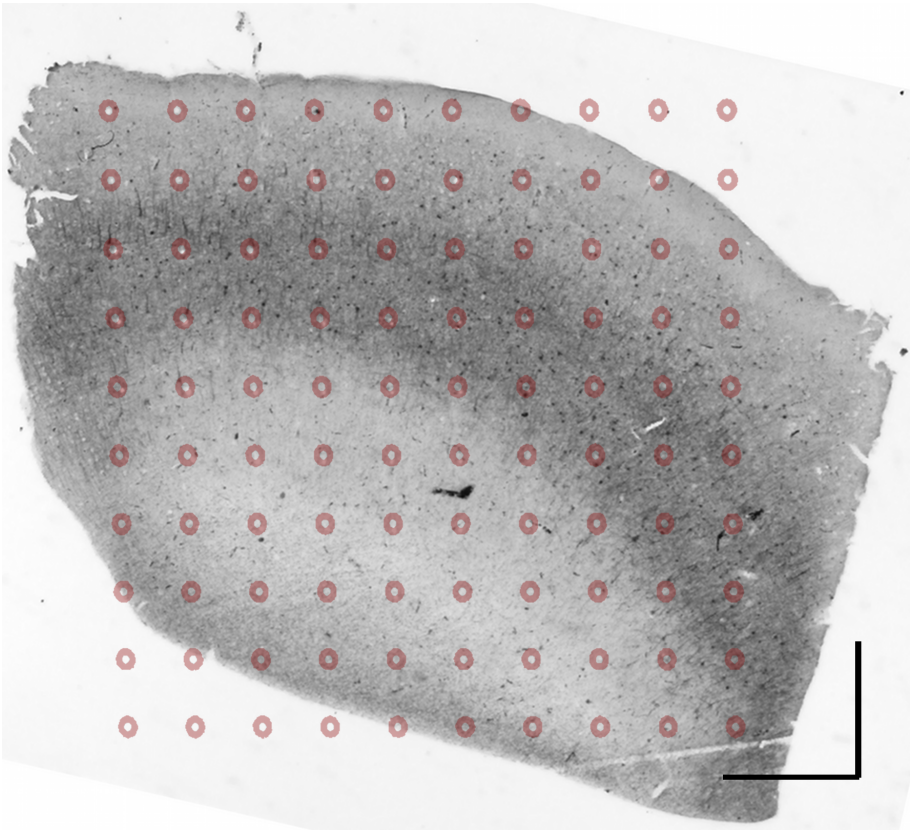

Figure 4.2: Example macaque neocortical slice, stained for non-phosphorylating neurofilament. Utah electrode positions are indicated with translucent red circles. Scale bars: 800 microns. Slice image provided by Matthew Ainsworth.

tion of the network. Any such model will inevitably only provide a caricature of the real biological system, but the increasing availability of experimental data has enabled modellers to improve and refine their caricatures. Discrepancies between experimental and modelling results can also be enlightening, potentially suggesting areas in which we need to refine our knowledge of the experimental system in order to create a more accurate model.

4.3.1 Model structure and connectivity

In order to create a model that incorporated the elements identified above, we took inspiration from published anatomical data to specify the different neuron types, their locations, and their connectivity. The model comprises fifteen neuron groups, defined in Table 4.1. It is designed to contain a similar number of neurons to a comparison experimental slice. This was calculated to be 175 421 neurons, based on the slice dimensions and neuron density of 38,335 neurons per cubic millimetre (Binzegger et al., 2004). The slice has clear spatial boundaries: neurons cannot be

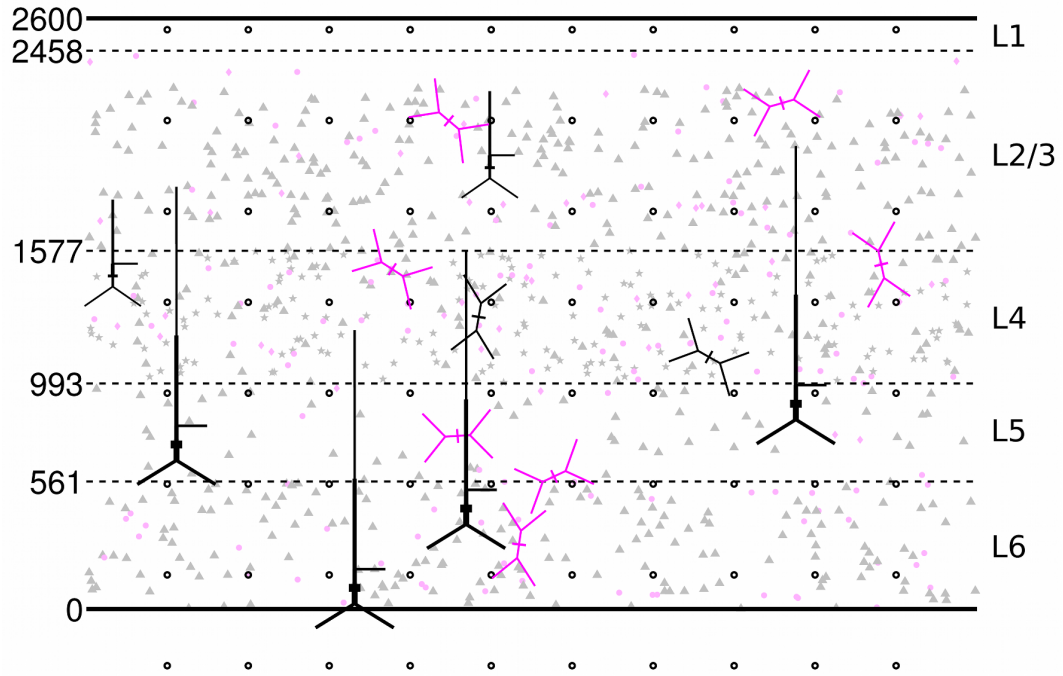

Figure 4.3: Slice model structure. Layer boundaries are given in microns. Subsets of soma locations from each neuron group are shown in faded black for excitatory neurons, or faded magenta for inhibitory neurons. Triangles represent pyramidal neuron somas, stars are spiny stellate cell somas, circles are basket interneuron somas and diamonds non-basket interneuron somas. One example full cell is shown for each neuron group, in solid black for excitatory neurons or solid magenta for inhibitory neurons. Compartment lengths are to scale; compartment diameters are not. Black circles are virtual electrode positions (first 8 rows shown).

Abbreviation	Neuron group description	Proportion of total model (%)	Compartments
P23	pyramidal neurons in layer 2/3 (L23)	27.4	8
SS4(L4)	spiny stellate neurons in L4 projecting to L4	9.7	7
SS4(L23)	spiny stellate neurons in L4 projecting to L23	9.7	7
P4	pyramidal neurons in L4	9.7	8
P5(L23)	pyramidal neurons in L5 projecting to L23	5.0	9
P5(L56)	pyramidal neurons in L5 projecting to L56	1.4	9
P6(L4)	pyramidal neurons in L6 projecting to L4	14.1	9
P6(L56)	pyramidal neurons in L6 projecting to L23	4.7	9
B#	basket interneurons in L#	13.7*	7
NB#	non-basket interneurons in L#	4.7*	7

*Proportions given for the whole model rather than per layer; proportions per layer are given in the Appendix

Table 4.1: Neuron groups, abbreviations, and number of compartments within our model. Basket interneuron groups are in L23, L4, L5 and L6. Non-basket interneuron groups are in L23, L4 and L5. Compartmental structures are shown in Chapter 2.

positioned outside of the slice edges, and axons cannot 'wrap around' these boundaries. We therefore required a connectivity model that would produce a suitable number of synapses given the large number of neurons, and that took into account each neuron's position in relation to the slice boundaries.

To set the number of connections between neurons of each type, we used the anatomical data from (Binzegger et al., 2004) (see Appendix). In this paper, the authors made reconstructions of 39 neurons from cat area 17, estimating connection numbers between groups using an improved version of Peters' rule (Peters and Payne, 1993). This provides one of the most complete quantitative descriptions of neocortical connectivity in the literature. Recently, Potjans and Diesmann (2012) created an integrated connectivity map using this anatomical data as well as physiological data from paired-pulse measurements (Thomson et al., 2002). We decided to use the numbers directly from (Binzegger et al., 2004) because the distinctions in connectivity between different interneuron sub-types are lost, while other aspects of the maps are quantitatively similar.

The data in (Binzegger et al., 2004) provide the number of neurons and connections, but not the spatial pattern of those connections. We specified a 2D Gaussian spatial profile to model the decay in connection probability with increasing distance from a presynaptic neuron, as data from layer 2/3 in rat visual cortex suggests that connection probabilities decay away from the presynaptic neuron with a this profile

	L1	L2/3	L4	L5	L6
P23	0.55	1.12	0.15	1.00	0.15
B23	-	0.50	0.15	0.15	-
NB23	0.20	0.20	0.20	0.20	0.20
SS4(L4)	-	0.30	1.12	0.40	0.15
SS4(L23)	0.15	0.40	0.50	0.15	0.15
P4	0.15	1.12	0.15	0.15	0.15
B4	-	-	0.50	0.55	-
NB4	0.20	0.20	0.20	-	0.20
P5(L23)	0.15	0.40	0.30	0.20	0.25
P5(L56)	-	-	0.15	0.50	1.00
B5	-	-	-	0.50	-
NB5	0.20	0.20	0.20	0.20	0.20
P6(L4)	-	0.15	1.00	0.15	0.15
P6(L56)	-	-	0.15	0.50	1.00
B6	-	-	-	0.10	0.50

Table 4.2: Axonal arborisation radii for each neuron group in each layer (mm), adapted from Figure 8 of the supporting information of (Izhikevich and Edelman, 2008). Where no radius was given for a neuron group in a layer in which connections are specified in the connectivity table given in the Appendix (marked with hyphens here), we set the radius to 0.05mm.

(Hellwig, 2000). The standard deviation parameter of the Gaussian profile was set using axonal arborisation radius measurements reported in (Blasdel et al., 1985; Fitzpatrick et al., 1985), as adapted in (Izhikevich and Edelman, 2008). These were different for each neuron group in each layer (see Table 4.2). Finally, we modelled the effect of slice cutting on connectivity by reducing the number of connections a presynaptic neuron could make by the proportion of the integral of its Gaussian connectivity profile that fell outside the slice boundaries, as described in the previous chapter.

Figure 4.4 shows the number of connections between neuron groups compared with the original numbers specified in (Binzegger et al., 2004). The proportional reduction in synapses is not the same for each connection type because of the varying axonal arbour radii. These reductions are important to consider when assessing the effect of connectivity changes on dynamics, but they illustrate that the general profile of connections between neuron groups is not substantially altered - connections from P23 to P23 and P5 neurons remain the most numerous, for example. Modelling thinner slices, or different axon arbour profiles, could lead to the over- or under-representation of particular connections in the model.

We use connectivity data from cat neocortex, but our comparison data was recor-

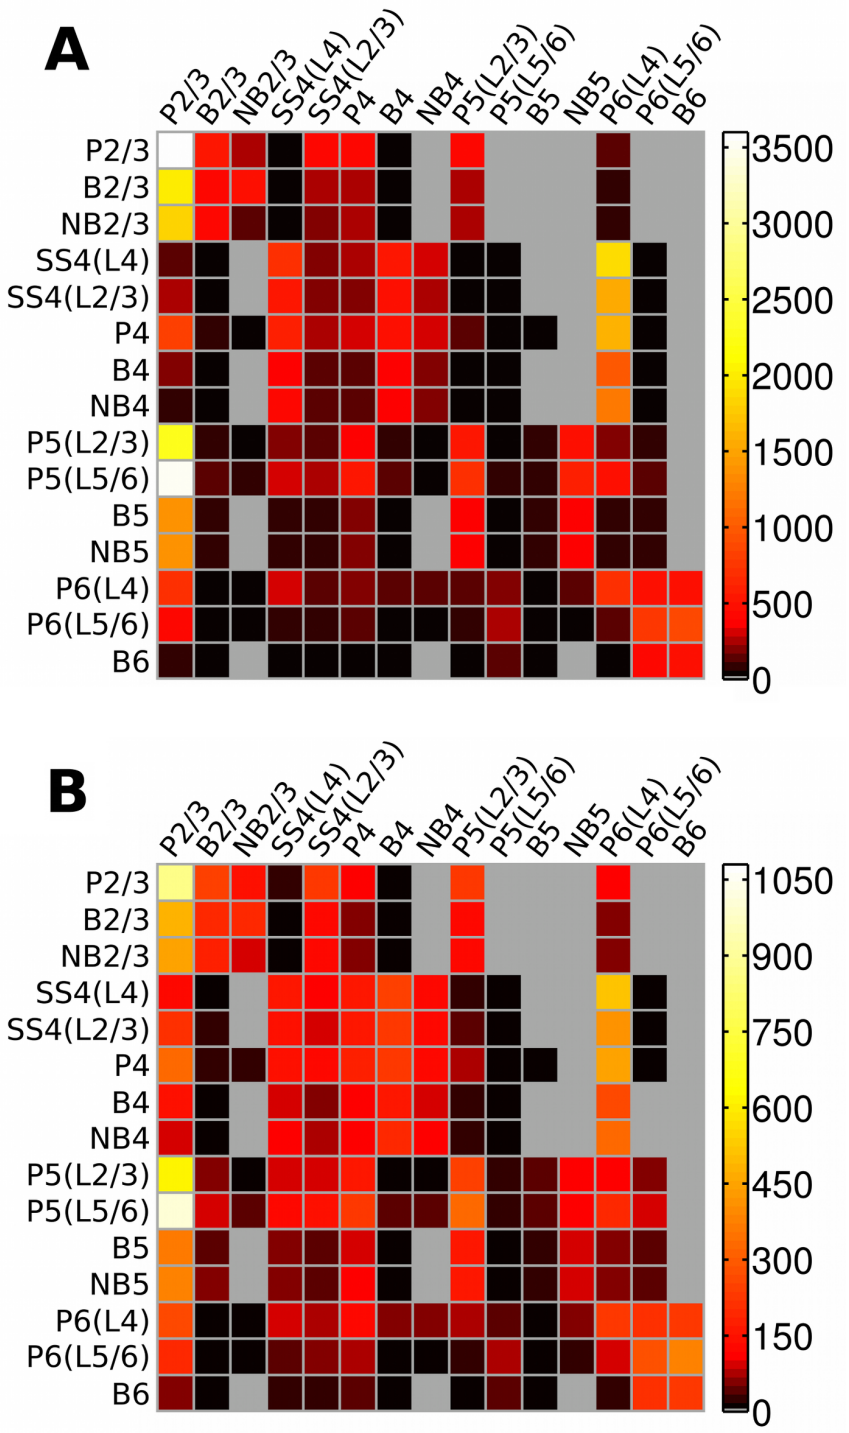

Figure 4.4: Changes in connectivity between neuron groups after slice cutting. (A) Expected number of connections from population of presynaptic neurons (columns) onto single postsynaptic neurons (rows) before slicing. (B) as (A), but after slice cutting. While overall connection number decreases (note different scale bars), some connections are affected more than others because of differing axonal arborisation sizes.

ded in macaque neocortex. The validity of the comparison between our model and the experimental data is therefore in question. Previous anatomical studies have emphasised the similarities between cat and macaque cortical connectivity (Callaway, 1998). Additionally, data collected by our experimental collaborators shows remarkable similarities in persistent gamma oscillations across different species, including rat, marmoset, macaque and human. Spike rates, peak frequency and spatial variation are all similar between species, suggesting that particularities of anatomy are not particularly relevant for the generation of gamma rhythms (Ainsworth, 2013).

4.3.2 Single neuron properties

Each neuron was represented using the default passive compartmental models with AdEx dynamics at the soma. As gamma oscillations can be generated by networks containing many different neuron types with many different individual neuron parameter values, we did not attempt to fit the spiking behaviour of each neuron type quantitatively to experimental data (in fact, (Brette and Gerstner, 2005) reported that the adaptation parameters of the AdEx model can be varied greatly and make little difference to the output spike trains produced during high conductance states). Instead we chose parameters to reproduce the general patterns of spiking behaviour observed in real neurons, using the membrane potential traces reported in Appendix A of (Traub et al., 2005b) as a guide ((Traub et al., 2005b) describes a layered cortical network model used to simulate, amongst other behaviours, persistent gamma oscillations *in vitro*). Each cell type's passive parameters were defined by its morphology and the electrotonic parameters given in Table 4.3; therefore, the parameters adjusted to fit the spiking responses of the Traub neurons were the spike slope factor Δ_t, threshold V_t, adaptation time constant τ_w, adaptation coupling parameter α, reset value v_{reset}, and instantaneous adaptation current increase β. We employed a qualitative approach to parameter adjustment, guided by the analysis of the AdEx model in (Naud et al., 2008). According to the classifications in (Naud et al., 2008), B, SS and P6 cells have a sharp reset, while NB, P23, P4 and P5 cells have a broad reset. B cells are non-adapting; SS and P6 cells are adapting; P23, P4, P5 and NB cells show an initial burst.

The model includes conductance-based alpha-Amino-3-hydroxy-5-methyl-4-isoxazolepropion
acid (AMPA, reversal potential 0mV) and $GABA_A$ (reversal potential -75mV) syn-

Neuron type	C_m (μFcm^{-2})	R_m (kΩcm^{-2})	R_a (Ωcm)	E_l (mV)	V_T (mV)	Δ_T (mV)	α (nS)	τ_w (ms)	β (pA)	v_{reset} (mV)
P23, P4	2.96	6.76	150	-70	-50	2.0	2.60	65	220	-60
SS4	2.95	5.12	150	-70	-50	2.2	0.35	150	40	-70
P5	2.95	6.78	150	-70	-52	2.0	10.00	75	345	-62
P6	2.95	6.78	150	-70	-50	2.0	0.35	160	60	-60
B	2.93	5.12	150	-70	-50	2.0	0.04	10	40	-65
NB	2.93	5.12	150	-70	-55	2.2	0.04	75	75	-62

Table 4.3: Neuron model parameters, used in all reported simulations (simulations of purely passive neurons only have C_m, R_m, R_a, and E_l specified.

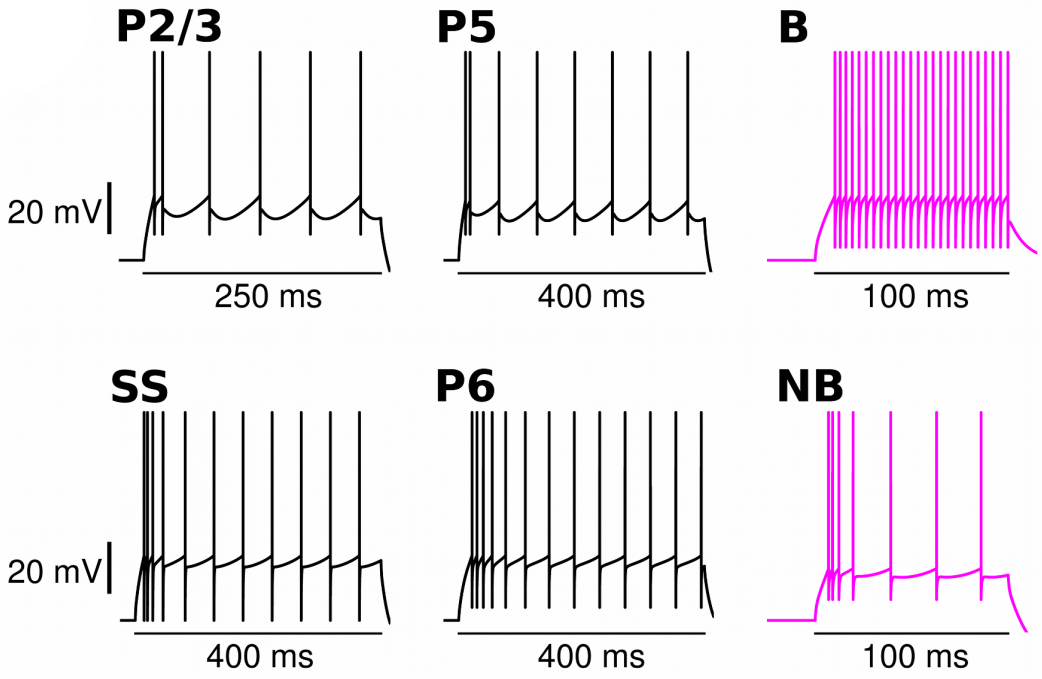

Figure 4.5: Responses to step-current injections into the soma compartment of each neuron type. Spikes are detected and cut-off at V_t + 5 mV; we extend the spike trace up to +10mV for illustrative purposes. Step current magnitudes were 0.5nA for the P23 neuron, 0.333nA for the SS neuron, 1.0nA for the P5 neuron, 0.75nA for the P6 neuron, and 0.4nA for the B and NB interneurons, as in (Traub et al., 2005b).

	E_{AMPA} (mV)	E_{GABA} (mV)	τ_{AMPA} (ms)	τ_{GABA} (ms)
P	0	-75	2.0	6.0
SS	0	-75	2.0	6.0
B	0	-75	0.8	3.0
NB	0	-75	0.8	3.0

Table 4.4: Synaptic parameters, taken from (Traub et al., 2005b). Row headers list the postsynaptic neuron type.

apses, the minimal set required for persistent gamma oscillations(Fisahn et al., 1998). Synaptic weights (Table 4.4) were chosen based on those reported in (Traub et al., 2005b), scaling the weights according to the number of synapses between groups in our model compared with the Traub model. Our neuron populations did not match theirs exactly, with the following differences (in addition to different numbers of neurons and synapses): our model includes interneurons in every layer, while the Traub model has only "superficial" and "deep" interneurons (with the deep interneurons providing inhibition to layer 4); the Traub model only has spiny stellate cells in layer 4 (no pyramidal or interneurons); the Traub model contains fast rhythmic bursting pyramidal cells in layer 2/3 and intrinsically bursting pyramidal cells in layer 5 - our model contains no bursting neurons; our model contains synapses between some neuron groups that are not present in the Traub model. We therefore had to make several arbitrary decisions when setting some synapse weights between groups. We give the synaptic weights in the Appendix.

We stimulated our model to mimic the bath application of kainate. This stimulates the pyramidal cell axonal plexus, providing the neurons with excitatory drive (see next chapter). We simulate this by applying independent random input currents to each neuron (Börgers and Kopell, 2005), modelled as Ornstein-Uhlenbeck processes as described above. The random current is distributed across the neuron's compartments proportionally to the compartment membrane areas.

4.3.3 Example model and outputs

In slice experiments with nanomolar kainate concentrations, the gamma oscillation is driven by L23, where neurons receive noisy excitatory drive from the excited axonal plexus of P23 neurons (Ainsworth et al., 2011; Cunningham et al., 2003, 2004b). We simulate this by providing a relatively large noisy current to P23 neurons, similar to (Ainsworth et al., 2011; Börgers et al., 2005). We set synaptic strengths [based on

	Mean current (pA)	Standard deviation (pA)	Noise correlation time constant (ms)
P23	360	110	2.0
SS4	205	50	2.0
P4	250	70	2.0
P5	860	260	2.0
P6	660	170	2.0
B	200	60	0.8
NB	160	40	0.8

Table 4.5: Random current input parameters.

(Traub et al., 2005b)] and noise currents to match the spiking activity and observed membrane potential fluctuation sizes reported in previous studies *in vitro*. Model parameters are given in tables the Appendix.

As described in previous experiments (Ainsworth et al., 2011; Cunningham et al., 2003, 2004b; Traub et al., 2005a,b), P23 neurons spike infrequently, while B23 neurons spike on most oscillation periods. Excitatory neurons in L4 do not take part in the oscillation (though still spike infrequently), while L4 interneurons are weakly entrained to the oscillation. In addition to the L23 gamma, the comparison slice exhibited increased gamma power in part of the infra-granular layers (see Figure 4.7 A, electrodes 6, 7, 16, 17, 26, 27), presumably caused by L5 as in (Ainsworth et al., 2011). We therefore used a relatively high coupling strength of P5 to B5 and NB5 neurons and a larger noisy drive current to L5 neurons to enable the L23 gamma to generate gamma in L5. The L5 gamma oscillation also weakly entrained L6 neurons to the oscillation.

The resulting spiking behaviour is shown in Figure 4.6, which shows a spike raster for 5% of the neurons in the model, along with example somatic membrane potential traces for each neuron group. The spike raster reveals that neurons near the slice x-boundaries (neurons nearest the cyan boundary markers in Figure 4.6) are less strongly entrained to the oscillation than neurons in the centre of the slice, because they receive fewer inhibitory inputs than more central neurons. This can be seen more clearly in the gamma power map in Figure 4.7. This effect is difficult to verify experimentally: recording spikes from enough neurons simultaneously is currently not possible, and the general patchiness of the LFP in experimental recordings means that it would not be possible to establish whether reductions in gamma power at

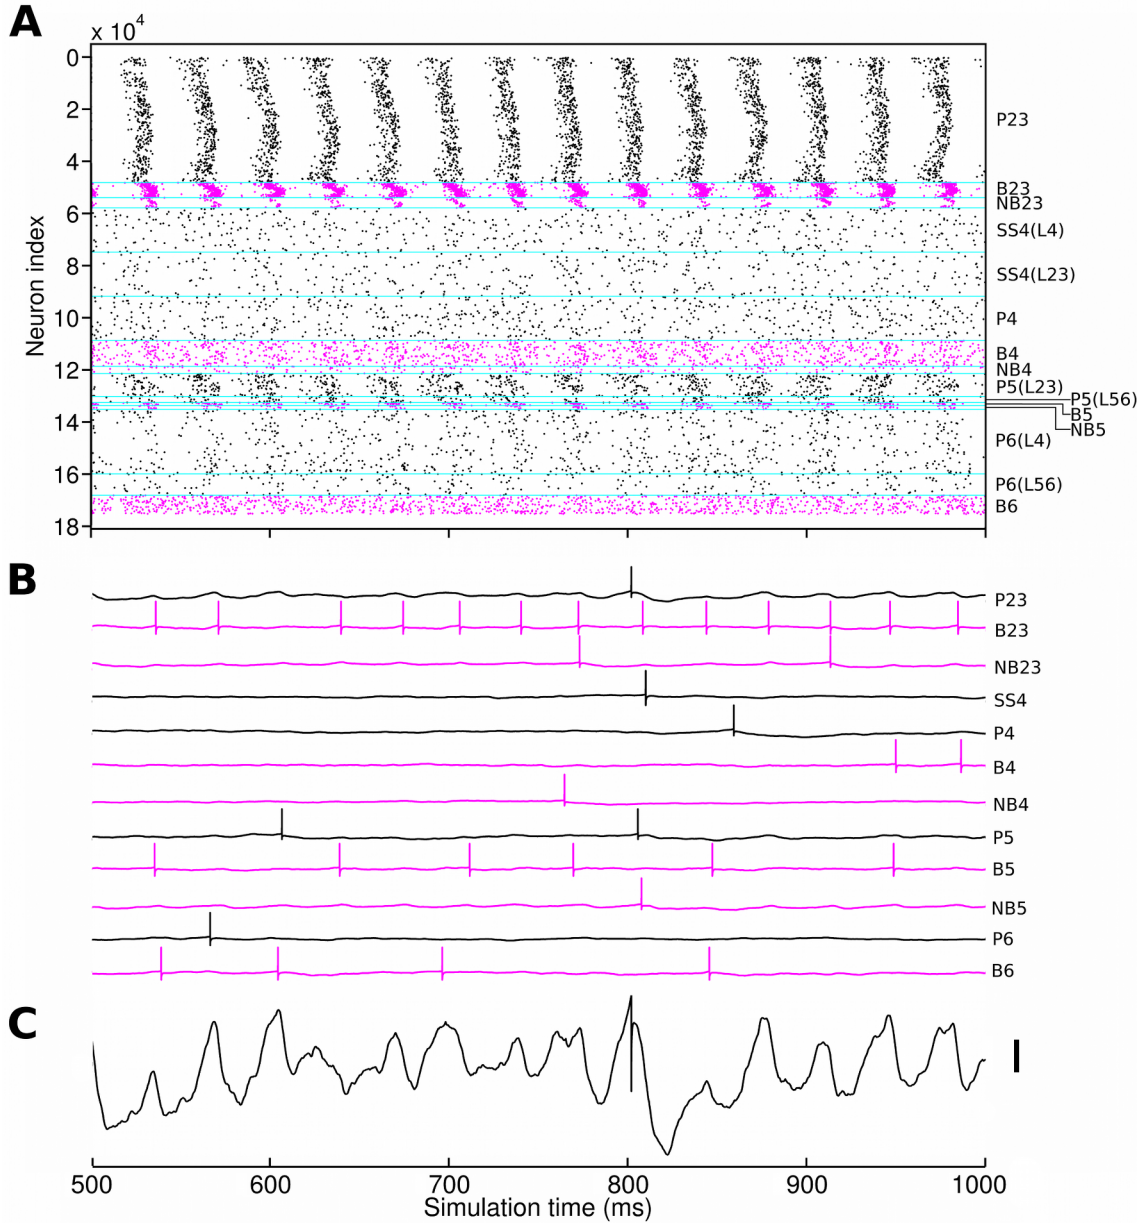

Figure 4.6: Spike raster and individual neuron responses during gamma oscillation. We use black to represent excitatory neurons and magenta for inhibitory neurons. (A) Spike raster showing spiking activity of 5% of all the neurons in the model (reduced number shown for clarity). Boundaries between neuron groups marked in cyan. Note strong persistent gamma oscillation in L23, with weaker oscillation in L5. Neuron IDs are arranged such that the lower the ID of the neuron (relative to the IDs of the neurons in its group), the further left on the x-axis the neuron is in the model. (B) Example soma membrane potential plots for the various neuron types. Most neurons fire sparsely, while B23 and B5 neurons fire on most oscillation periods. Spikes are cut-off at V t + 5 mV in the simulation; we extend them up to 10 mV here for illustrative purposes. (C) Close-up of P23 neuron soma membrane potential (cut-off -45 mV). Scale-bar: 5 mV.

slice edges were a result of slice cutting or other tissue properties that alter the gamma oscillation power. However, the effect of slicing on the dynamics at the slice edge observed here should be taken into account when interpreting the LFP recorded near slice boundaries *in vitro*.

4.3.4 Comparison with experimental data

Having verified that the model produced the expected spiking output, we looked at the simulated LFPs and compared them with those recorded *in vitro*. The raw experimental data of MEA recordings from resected macaque temporal cortex was kindly provided by Matthew Ainsworth at the Newcastle University Institute of Neuroscience (now at the University of York). Persistent gamma oscillations were induced by bath application of kainate (400nM - 800 nM). The recordings were first re-referenced to the common average, then resampled at 1 kHz (Matlab Signal Processing Toolbox resample() function, which applies an anti-aliasing filter). We removed line-noise and harmonics by band-pass filtering each recording at 49-51 Hz, 99-101 Hz, 149-151 Hz, 199-201 Hz and 249-251 Hz (symmetrical Butterworth filter, 8th order) and subtracting the resulting signal from the original signal. The recordings were then band-pass filtered between 2 Hz and 300 Hz (symmetrical FIR filter, Kaiser window, 2000th order). Filters were run forwards and backwards so no phase distortion was introduced (reported filter orders take this into account). We restricted our analysis to an 18 second segment of the recording that was identified as artefact-free in all channels by visual inspection of the filtered traces. After filtering, these segments were normalised to zero mean, unit standard deviation to facilitate signal comparison across the MEA. Total LFP power at each electrode varied considerably (possibly due to variable electrode impedances in the tissue, and possibly because of other non-gamma activity that we were unaware of) and as we were interested in the gamma band as it related to rest of the signal, this allowed us to account for this overall signal level variation when considering gamma power over the whole electrode array. As the total signal power was not nearly so variable in the model data - electrode impedances are constant and no activity other than the gamma oscillation is present - we did not normalise the LFP recordings from the model in this way. As the model LFPs also contained prominent gamma harmonics that varied in power depending on electrode position (see Figure 4.9),

and which were absent in the experimental recordings, this kind of normalisation would not have been suitable for comparing gamma power across space in the model recordings.

Figure 4.7 shows a comparison over the whole electrode array between the model and the experimental recordings. Figure 4.7 A shows a power map over the MEA of the signal magnitude in the low-gamma frequency band, calculated by integrating the power spectrum of the LFP at each electrode (calculated using the Thomson multitaper method - Matlab Signal Processing Toolbox pmtm() function) between 20Hz and 40Hz and using bicubic interpolation (Matlab interp2() function) to estimate the power between electrodes. The shape of the experimental neocortical slice is apparent as electrodes placed in the white matter (bottom left) or above the cortical surface (top right) are greyed out. As predicted by previous research, strong gamma power is apparent in the supra-granular layers. The gamma power at each electrode is highly variable, resulting in a patchy power map. This is not captured by the model, whose structure is homogeneous along the x-axis (the power map was constructed in the same way as for the experimental recordings, using 1.5s simulated data and without common average re-referencing or normalisation - see above). To check that the gamma oscillation profile across the slice was in fact "patchy" and that the gaps between patches were not simply a result of fluctuations in the power of an oscillation being generated by the same network across space, we found the cross-correlations between the LFPs from 9 electrodes in the two largest gamma patches shown in Figure 4.7 A and all other electrodes in the grid (again, excluding those electrodes in the white matter or above the slice). The results are shown in Figure 4.8. To create this figure, we first band-pass filtered the experimental LFPs between 25 Hz and 40 Hz. We then took the cross-correlations between the LFP at the selected electrode in layer 2/3 and the LFPs at all other electrodes, with a maximum lag time of 100 ms. We found the maximum (absolute) correlation values within this lag time for each electrode comparison. The decay in correlation strength was similar for the nine electrodes we chose to investigate, and no correlation structure matching the gamma power over space emerged. We took this to suggest that the gamma oscillation was indeed "patchy": the oscillation measured at each electrode (or at close-by electrodes) is likely to be generated by subsets of neurons that are not taking part in the same collective oscillation as the

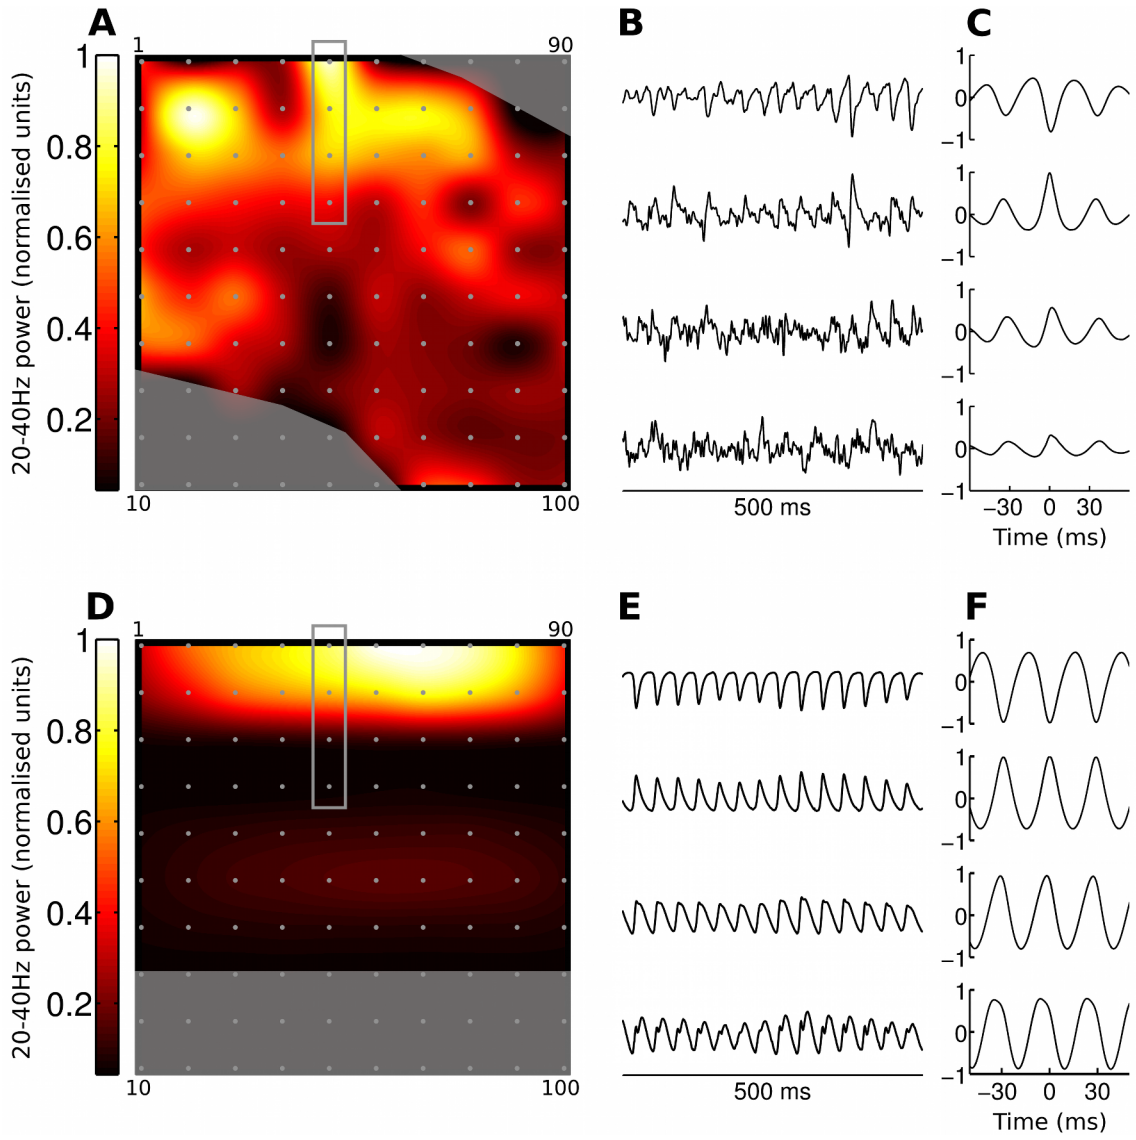

Figure 4.7: Comparison of experimental and simulated MEA recordings. (A) Map of gamma frequency power across the electrode array. Electrode positions shown as grey dots, corner numbers indicate electrode IDs. Shaded areas show where electrodes were discounted because they fell either outside the slice boundaries or within the white matter. Gamma power is strongest at the top of the slice, corresponding to L23. (B) Example LFP traces from electrodes 41 to 44 [indicated by grey rectangle in (A)]. Traces have been normalised to unit standard deviation for ease of comparison. (C) Cross correlation of signals from electrodes 41 to 44 with signal from electrode 42, illustrating phase inversion in the signal from electrode 41. This electrode was identified as being in layer 1 by post hoc histology (not shown). Gamma map & cross-correlations estimated from 18s of data. (D-F) As (A-C), but for the neocortical slice model (gamma map & cross-correlations estimated from 1.5s of data). The model produces a strong gamma oscillation in L23, with a weaker oscillation in L5. LFPs in E have been normalised to unit standard deviation for ease of comparison, but non-normalised LFPs were used to calculate the gamma map in D (see main text).

neurons at more distant electrodes.

The phase inversion between layer 1 and layer 2/3, illustrated in Figure 4.7 B and C, does emerge in the model (Figure 4.7 E and F). This is from the positioning of current sinks and sources on the P23 neurons during the gamma oscillation. This result is in agreement with the source-sink interaction mechanism of phase inversion demonstrated experimentally in kainate-induced gamma oscillations in entorhinal cortex *in vitro* (Cunningham et al., 2003) and which we explore further in the next chapter. The cross-correlations between electrodes shown in Figure 4.7 C and F also reveal how the strong gamma oscillation in L23 dominates across the electrodes more than in the experimental recordings. This is, again, a result of the synchronous activity in L23 along the x-axis in the model, meaning that the LFP signal created by the gamma oscillation is not degraded by influences from the non-oscillating areas in the slice as occurs *in vitro*. Finally, figure Figure 4.9 shows large peaks at harmonics of the base gamma frequency. These peaks result from the shape of the oscillatory signal, which is closer to a saw-tooth wave than a sine wave. This shape is very clear in the simulated LFP (figure Figure 4.7 E), as the coherent layer 2/3 oscillation totally dominates the signal. By contrast, in the experimental recordings, the gamma frequency harmonics in the power spectrum are mostly concealed by noise except at the electrodes where gamma power is strongest, where the first gamma oscillation harmonic is visible in the power spectrum (see Figure 5.17 in Chapter 5, which shows LFP power spectra at electrodes 41 and 42 as indexed in Figure 4.7). In summary, our model appears to capture the LFP features observed during gamma oscillations through the depth of the slice, but not laterally, a discrepancy that could arise through lateral inhomogeneities in connectivity patters, synaptic strengths, or cell densities. Further models and experiments should investigate these possibilities.

4.3.5 Checking the influence of the AdEx spiking mechanism

The LFPs recorded from the model contain potentially unrealistic contributions from the AdEx spiking mechanism in the model neurons (see above). To check if these contributions significantly affected the results, we reran the simulation with the AdEx mechanism in each neuron removed, using VERTEX's spike-import feature to replay the spiking activity in the network. The results for two of the electrodes in the virtual MEA are shown in Figure 4.9. In this model, with a large number of

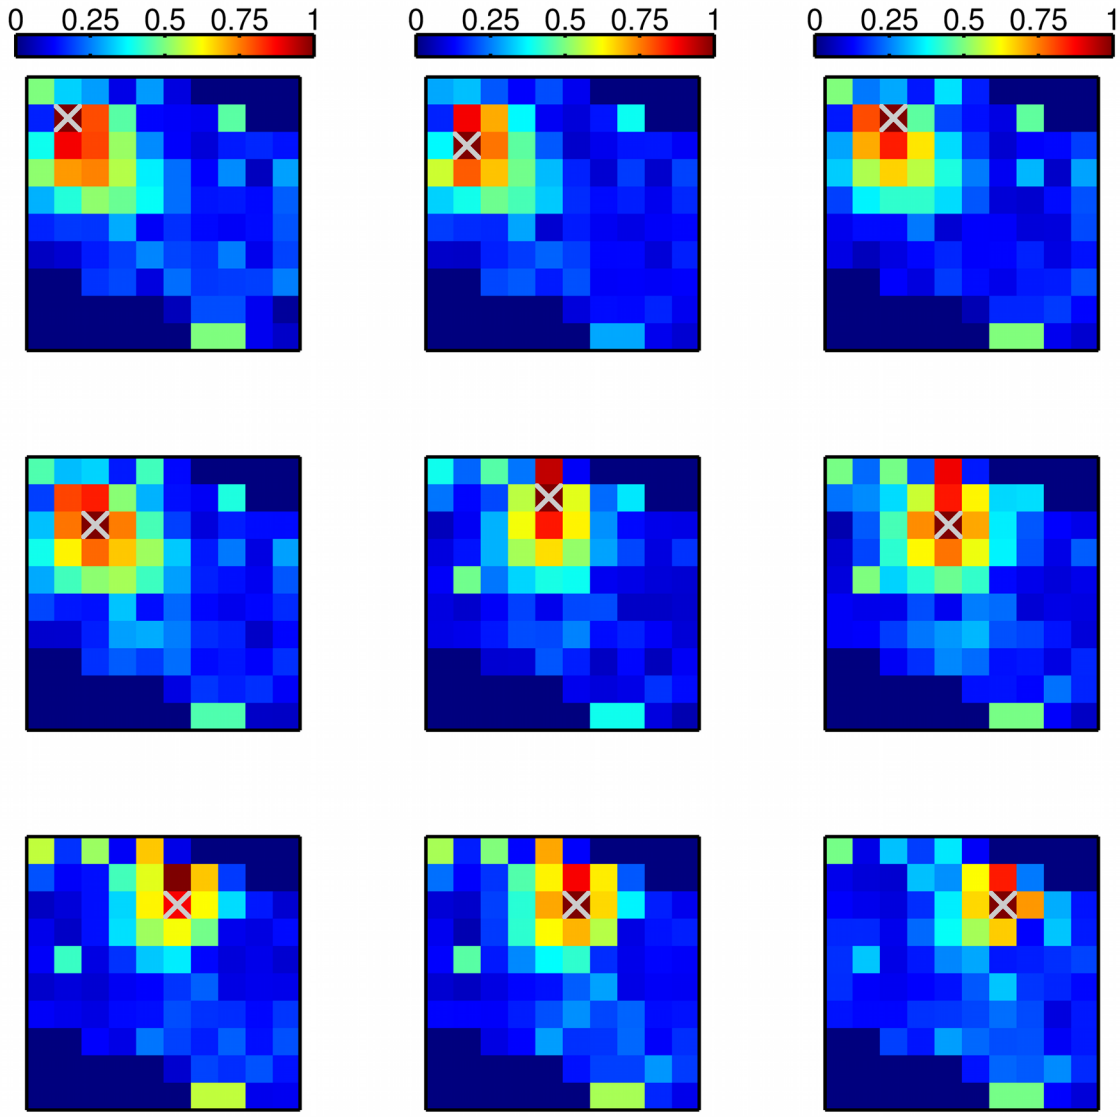

Figure 4.8: Correlations between 9 electrodes in the gamma patches and all other electrodes in the slice. The reference electrodes for the cross-correlations are marked with a cross. All reference electrode positions are in the layer 2/3 gamma patches revealed in Figure 4.7 A. Absolute values of negative correlations were taken to produce the correlation maps. The decay in correlation strength around each selected electrode is quite similar, with no long-range correlations apparent. Specifically, a correlation structure linking the different gamma patches is not apparent (compare with Figure 4.7 A), suggesting that the gamma power is indeed "patchy" and that separate local circuits could be generating the gamma oscillation at different points in the network. The comparison electrodes (from top left, moving columnwise) are 12, 13, 22, 23, 42, 43, 53, 63, 73.

neurons receiving synchronous synaptic input, the influence of the AdEx mechanism compared with the synaptic input is very small. Electrode 97 has presumably been positioned close to the soma compartment of a neuron, and the spike in the extracellular potential reflects the instantaneous membrane potential reset of the AdEx mechanism. Many of these very rapid events will have been missed during the simulation, as the extracellular potential is only sampled at 1kHz, while the time-step of the simulation is 0.0325ms. The spike causes significantly increased power above ∼100Hz compared with the extracellular potential recorded from the purely passive model with spike replay, but the spectrum is unaffected at lower frequencies. We therefore suggest that, though analysing simulated LFPs from models containing multi-compartment AdEx neurons can be useful during explorative model runs, they should be checked against LFPs produced by an equivalent passive model using the spike-replay functionality. This is particularly important when the LFP contributions from the summed synaptic currents do not dominate the influence of the AdEx mechanism, such as in smaller models with less synchronous firing.

4.4 Summary

We created a large-scale, spatially organised model of a neocortical slice to use in conjunction with MEA recordings *in vitro*. The model is similar in scale to the brain slice preparation - an experimental method that allows the study of isolated network dynamics, and that has proved extremely useful in elucidating the mechanisms of various patterns of network activity (Traub and Whittington, 2010). As we are interested in how the LFP relates to network activity across space, simulations containing a number of neurons approaching an experimental preparation are preferable so that the LFP is not distorted by overly strong synaptic currents (as frequently used in reduced network models so that neurons receive adequate input to fire), and each neuron contributes an appropriate amount to the simulated LFP signal. A model on this scale also allows a direct comparison to be made with the experimental recordings.

However, several problems remain when investigating a model of this scale and incorporating as much detail. Firstly, a large number of parameters must be estimated. Experimental measurements can be used as constraints, but these are often

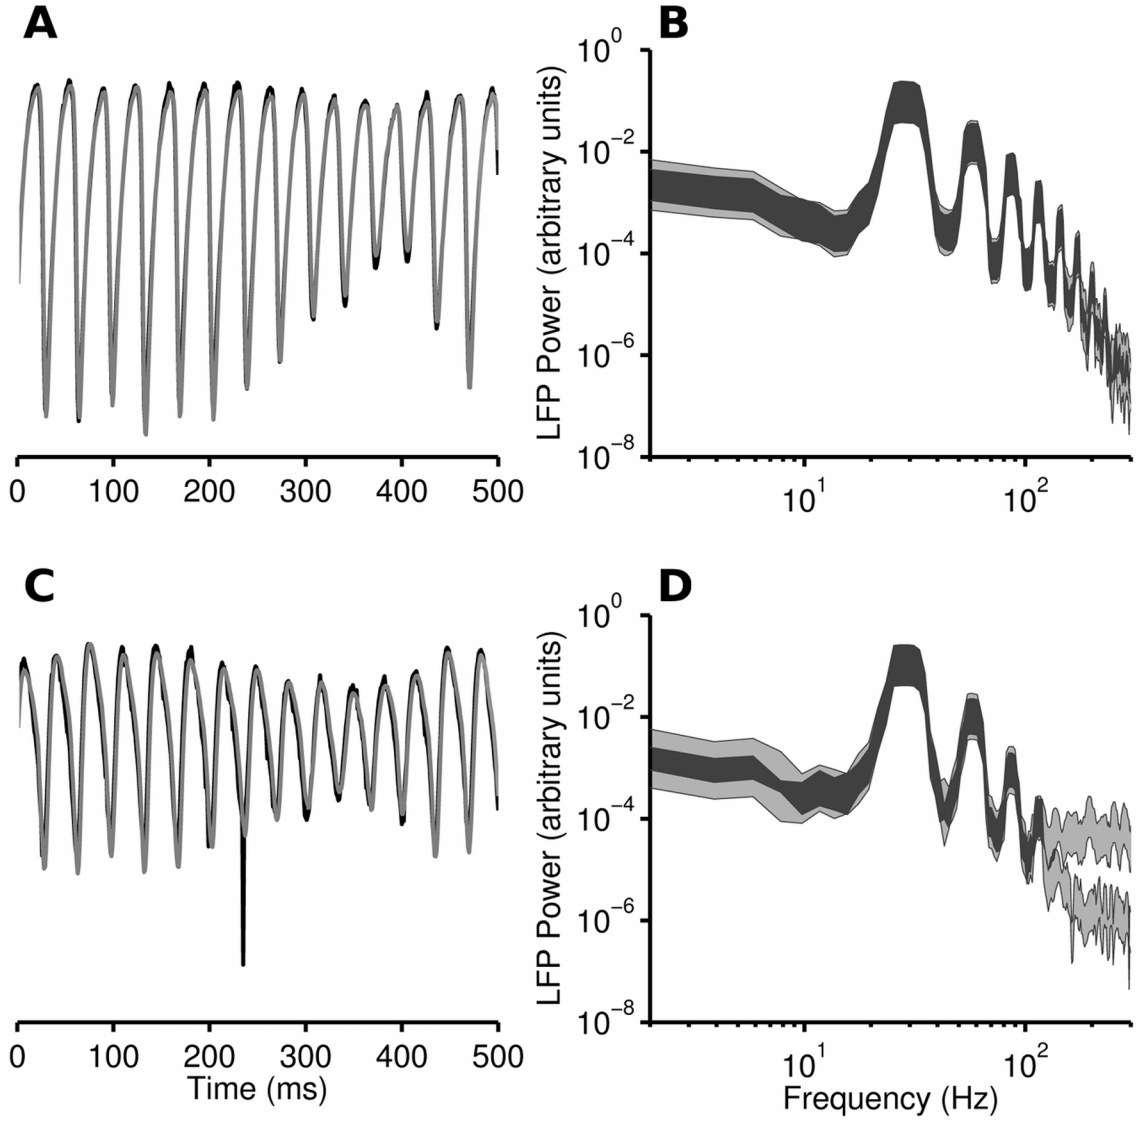

Figure 4.9: Comparison of simulated LFPs when using purely passive neurons with imported spike times, and when using the AdEx spiking model in the neocortical slice model. Traces in (A) and (C) have been normalised to zero mean, unit standard deviation so that the shape of the LFP signal can be more easily compared, but are otherwise unfiltered. The adaptive current of the AdEx mechanism introduces an offset in the simulated LFPs, but does not dramatically affect the shape of the signal. (A) Simulated LFPs from electrode 42; AdEx version in black, passive version in grey. (B) Power spectral density overlap of these signals (estimated for 500 ms signal), with overlapping parts of the estimated spectrum shown in dark grey and non-overlapping shown in light grey. (C) As A, but for electrode 97. Note the small, sharp spike at ∼230 ms in the AdEx signal. This is a result of the AdEx reset mechanism creating a very short, fast current from a neuron very close to the electrode. (D) As B, but for electrode 97. Power spectra diverge above ∼100 Hz due to high frequency spike contamination in the AdEx model, but match closely below 100 Hz.

statistical estimates, come from various different species, or are otherwise incomplete. Synaptic strengths are not well defined and can vary dramatically depending on activity and environment, and statistical connectivity estimates hide possible inhomogeneities in synapse density across space. We are therefore forced to make many arbitrary assumptions based on best guesses. Secondly, the analysis of the model must necessarily be more constrained than for a smaller, simpler model, as simulation and analysis take longer and visualisation becomes more challenging. Finally, while it is possible to compare simulation output directly to experimental recordings, various aspects of the model that cannot always be established from the comparison data influence the LFP. In our case, the experimental data used to create the model come from a different part of the neocortex and different species than the comparison experimental LFP recordings. While the relative uniformity of the neocortical circuit, and indeed persistent gamma oscillations, across species (Ainsworth, 2013) means that these details are not necessarily crucial when modelling the spiking statistics in the network, different neuronal densities, connection densities and even layer boundaries all have an effect on the measured LFP.

For the next step in our investigations, we decided to focus on the LFP during gamma oscillations in a simplified model of layer 2/3, allowing us to perform a more thorough analysis of particular features of the LFP during this activity regime. We describe this analysis in the next chapter.

Chapter 5

LFP features during gamma oscillations

5.1 Introduction

The main reason for our investigations into the LFP is so that experimental LFP recordings can be more easily interpreted in terms of the underlying neuronal activity. In the previous chapter, we described a large-scale model of persistent gamma oscillations in a neocortical slice, implemented with the VERTEX simulation tool. The virtual electrode functionality in VERTEX allowed us to simulate LFPs during this activity regime, so we could make a direct comparison with experimental extracellular recordings. This model reproduced certain features of the experimentally recorded LFP, but the spiking activity was too synchronous across the width of the slice model and the spatial LFP properties did not match comparison experimental recordings made using Utah MEAs.

To investigate the LFP during persistent gamma more thoroughly, we decided to focus only on the generating layer, creating simplified network models of layer 2/3. We imposed artificial conditions on the spiking activity in the models, so that we could systematically investigate LFP features across space due to precisely known and constrained spiking activity and connectivity. We used the set of models to investigate how the firing of each neuron type influenced the LFP, how layer depth affects LFP power and the location of the experimentally observed phase inversion, the visibility of the gamma rhythm outside of the network region that generated it, and the LFP's frequency scaling characteristics. We relate our results to various

experimental findings regarding the origin and features of LFPs.

5.2 Model of persistent gamma in neocortical layer 2/3

We chose to continue to base our gamma oscillation model on the *in vitro* slice preparation with bath-applied kainate, both because we had previous experience in modelling this particular type of gamma oscillation, and because the neuronal firing patterns (sparse pyramidal firing and relatively regular basket interneuron firing) appear to correspond with the *in vivo* case (Csicsvari et al., 2003; Hájos et al., 2004). Our results therefore have implications for synaptically-driven oscillations in general, rather than specifically *in vitro* gamma oscillations after kainate application. To simplify our analysis we concentrate only on layer 2/3, as this layer generates the gamma oscillation *in vitro*.

We chose to investigate three aspects of the LFP that we thought could help in the interpretation of experimental LFP recordings: the contribution of each neuron type to the total LFP, the spatial profile of the gamma power in the LFP, and the location of the phase inversion. Taking inspiration from the experimental observation of patchy gamma power, we decided to investigate these aspects in a model containing a gamma-generating region surrounded by regions that did not participate in the gamma oscillation (i.e. the neurons in these regions fired randomly). These enforced conditions allowed us to study the spatial profile of the LFP from differently sized gamma patches, while minimising the number of variables we changed between simulations. Additionally, we could make observations about properties of the LFP during random activity in the layer 2/3 network.

In the next sections we describe the composition, topology, connectivity and dynamics of the model, as well as the LFP calculation method for investigating different sizes of gamma generating networks embedded in surrounding tissue with randomly spiking neurons. We refer to a single model; in fact, we generated ten models with identical parameters but different random number seeds, so that the random connectivity and input currents in each model were different. We measured the reported quantities in each model realisation, and report the mean of these measurements in the results section. This was so that random effects in network

initialisation were averaged out, to ensure that the measurements we report are generic and not specific to a particular network realisation.

5.2.1 Model composition

The simplest network models of persistent gamma include only basket interneurons and pyramidal cells (Börgers and Kopell, 2003, 2005). However, other types of interneuron also fire during persistent gamma, but at a lower rate than the basket cells (Traub et al., 2005a). While they do not contribute to gamma generation, they affect the LFP as they primarily target pyramidal cell dendrites, so the resulting current sources and sinks have a different spatial profile from those created by basket cell synapses. The characterisation of the different interneuron types into soma-targetting or dendrite-targeting is extremely simplified, as a huge number of different interneuron categories has been identified (DeFelipe, 1997; DeFelipe et al., 2013; Jonas and Buzsaki, 2007; Kawaguchi and Kubota, 1997; Markram et al., 2004; Petilla Interneuron Nomenclature Group et al., 2008). We decided to use this characterisation - as in (Traub et al., 2005a,b) - because different roles for the different interneuron sub-types in persistent gamma have not been elucidated. We therefore included three cell types in our simplified layer 2/3 model: pyramidal neurons (excluding chattering cells), basket interneurons, and dendrite-targeting interneurons (which we refer to as non-basket interneurons, as before). We again abbreviate these populations as P, B and NB cells, respectively, and use A to refer to the three populations combined.

As in the previous chapter, we made use of the population size estimates in (Binzegger et al., 2004) to set the size of each population. We limited ourselves to modelling the *in vitro* slice, so specified a model width of 4 mm and a slice thickness of 0.4 mm. Neurons were positioned with their somas contained within a soma-layer of a certain depth, with the pyramidal neuron dendrites extending above the soma-layer towards an imaginary cortical surface. We wanted to investigate the effect of varying the soma-layer depth on the LFP while minimising the number of variables between conditions, so we chose the number of neurons given the shallowest soma-layer depth (100 microns) and kept this number constant. The neuronal density in layer 2/3 in the Binzegger et al. (2004) data is $\sim$69,000 mm^{-3}, resulting in a total model size of 11,000 neurons. There were 9000 pyramidal cells, 1000 basket cells

and 1000 non-basket cells, which we refer to as P, B and NB, respectively (as in Chapter 4). Neurons were initialised with their somas at random points within the soma-layer and slice boundaries, but their dendrites were allowed to extend outside them, so the top of the pyramidal neurons' apical tuft compartments could reach a maximum model height of 430 microns with a soma-layer depth of 100 microns. We also tested model depths of 200 microns and 300 microns, for which we kept the connectivity, x- and y-coordinates the same as the original 100 micron soma depth model but shifted each neuron along the z-axis by a constant multiplier of 2 or 3.

5.2.2 Connectivity

Connectivity between neuronal groups was calculated using the same method as in Chapter 4: connection numbers were taken from (Binzegger et al., 2004), and the connection probability decayed with a 2D Gaussian profile (Hellwig, 2000), with connection numbers adjusted to account for slice boundaries using the method described in Chapter 3. This introduces some heterogeneity into the number of connections each neuron makes: each cell's location in relation to the model boundaries is different, so the range of the integral for calculating the remaining number of connections for each cell will be different. We also reduced the number of connections made by each neuron by a half because of the reduced number of neurons in the model: we wanted to approximate the right level of network synchrony given the number of neurons, as this dramatically affects the LFP magnitude. The number of connections is consistent with the estimated convergence given a model size of 11,000 neurons in (Potjans and Diesmann, 2012), taking into account the extra missing synapses from slice cutting. We set the axonal arborisation radius to be 500 microns (i.e. 250 micron standard deviation parameter) for all neuron types so that the impact of this somewhat uncertain estimate on the LFP's spatial profile was the same for each neuron group. In fact, a uniform Gaussian distribution with a large radius, as used in the full slice model, may be a poor model for layer 2/3 pyramidal neuron connectivity in cat and macaque (at least), as it appears that these neurons make long-range projections to specific patches of postsynaptic neurons (Bauer et al., 2012; Binzegger et al., 2007; Kisvárday and Eysel, 1992) rather than targetting all neurons at the same distance with equal probability (see discussion in Chapter 4). The current simplified model, then, can be considered as ignoring these long-range

(A) Convergent connectivity

	P	B	NB	Total
P	694	97	61	852
B	440	81	86	607
NB	405	81	36	522

(B) Divergent connectivity

	P	B	NB
P	694	872	549
B	49	81	86
NB	45	81	36
Total	788	1034	671

Table 5.1: Simplified layer 2/3 model connectivity. (A) Average number of connections received by one postsynaptic neuron (rows) from each presynaptic group (columns): convergent connectivity. (B) Average number of connections made by one presynaptic neuron (columns) onto each postsynaptic group (rows): divergent connectivity.

connections and only considering local connectivity, for which a Gaussian profile appears reasonable (Hellwig, 2000). Modelling patchy projections would introduce too many variables into the present study, but future investigations should certainly investigate how the properties of the superficial patch system affect the LFP. Finally, unlike in the full slice model, we set the maximum distance for connections to 500 microns so that all connections are made within a known range. We kept the connectivity constant for each dynamics simulation performed. Mean connection numbers between groups are given in Table 5.1. Neurons could form synapses with the postsynaptic compartment IDs specified in Appendix 1.

5.2.3 Neuron dynamics

We use the reduced compartmental models described in Chapter 2 for the neuron morphologies, with interneurons having the same morphology as the reduced spiny stellate cell models as in the previous chapter. Using a different cell type morphology to represent smooth interneurons means that we cannot guarantee that the membrane currents of the interneuron populations contribute suitably to the LFP. However, we made the assumption that this would not affect the LFP greatly: pyramidal neurons are larger than interneurons and make up over 80% of the neurons in the model, so we assumed that the majority of the LFP signal is due to pyramidal cell membrane currents. We check this assumption below. Each neuron had an AdEx mechanism at the soma compartment to generate spikes, with parameters the same as for the equivalent layer 2/3 neurons described in the Chapter 4.

We used current-based synapses to maintain linearity (Lindén et al., 2010). This allowed us to consider each presynaptic population's contribution to the LFP in-

	P	B	NB
P	1.0	-3.5	-1.8
B	37.0	-1.0	-0.1
NB	37.0	-6.0	-0.2

Table 5.2: Synaptic weights for the simplified layer 2/3 model. These values represent the instantaneous increase in synaptic current, in pA, at a postsynaptic compartment (rows) on receiving a spike from a presynaptic neuron of the specified type (columns).

dependently by running the simulation with the imported spike times from one presynaptic group at a time. The total LFP is found by summing each presynaptic group's contribution (we used the same principal to calculate each postsynaptic group's contribution to the LFP to check our assumption about the contribution of pyramidal cells). We used single-exponential type synapses, with synaptic weights given in Table 5.2. To set these, we took the weights of the conductance-based synapses used in (Cunningham et al., 2004b), scaled them according to the number of connections in our model, and calculated the instantaneous current change due to these conductances given a membrane potential of -55 mV.

5.2.4 Network dynamics

The gamma oscillation was generated by applying independent randomly fluctuating excitatory currents to each neuron, as in Chapter 4. A single random current was spread across a neuron's compartments, scaled according to the compartments' relative membrane areas. While this scenario is unrealistic given the source of the noisy drive in persistent gamma, it allowed us to ignore the details of the stimulation source for the purposes of this investigation so we could focus on the LFP of the resulting network activity. A persistent gamma type oscillation emerged in this network. We ensured that the size of inhibitory postsynaptic potentials in the pyramidal cells matched those reported in (Ainsworth et al., 2011).

We also ran simulations in which neurons fired randomly, to simulate the network LFP generated under random spiking conditions. For these simulations, we did not apply any random currents and removed the AdEx mechanism from the somatic compartments. Each neuron's spike times were generated by an independent Poisson process. We simulated two randomly spiking regimes: a high activity regime with neurons firing at 5 Hz, and a low activity regime with neurons firing at 1 Hz.

Synaptic currents were integrated as normal, but did not influence cell spiking at all. Unless otherwise mentioned, the results reported here are with the high activity surrounding regions.

5.2.5 LFP simulation

We used various combinations of the gamma and random activity regimes to study the spatial properties of the LFP during localised gamma oscillations. In order to create a compound LFP, we added the LFP generated by the spiking of a central subset of neurons during gamma (we refer to this sub-population as P_γ, in gamma-generating region R_γ with spatial width X_γ) to the LFP generated by the spiking of the remaining surrounding neurons during random firing (we refer to this sub-population as P_R, in randomly spiking region R_R with spatial width X_R). This method allowed us to vary the size of the gamma generating patch of tissue, while keeping all other aspects of the simulation constant. Before creating this compound LFP signal, we removed any LFP contributions from the AdEx mechanism and input currents in sub-population P_γ by using the spike-loading function of VERTEX to replay the gamma spiking activity in the same network of purely passive neurons with no input currents.

We were interested in how far away from R_γ the gamma oscillation would be visible in the LFP, as well as the depth of any phase inversion both inside and outside R_γ, so we calculated LFPs at 1215 locations forming a 81 by 15 grid on the x-z plane. The grid was placed so that the bottom left electrode was at (0, -50) microns and the top right electrode was at (4000, 650) microns on the x-z plane, with a spacing of 50 microns and a constant y-value of 200 microns (see Figure 5.1). All reported spatial LFP values are therefore constrained by this grid placement and resolution.

For each experimental condition, we calculated the contribution of each population to the LFP separately, both from a presynaptic and postsynaptic perspective. The LFP for the whole model was found by summing these group contributions.

5.2.6 Definitions of measures

We quantify various features of the LFP using the following measures:

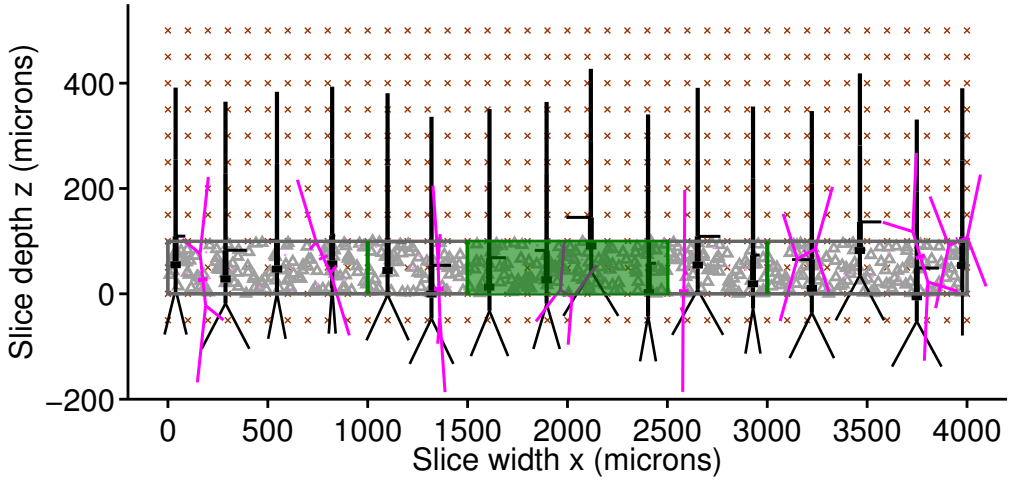

Figure 5.1: Experimental setup - slice anatomy. Shows 5% of cell soma positions within slice boundaries (grey) for 100 micron deep slice, plus examples of dendrite morphologies for pyramidal (black) and basket (magenta) cells. Electrode positions are shown as crosses. The area shaded green shows the region of cells participating in a persistent gamma oscillation. The vertical green marks to the left and right of this region mark the maximum range of synaptic connections made from this region.

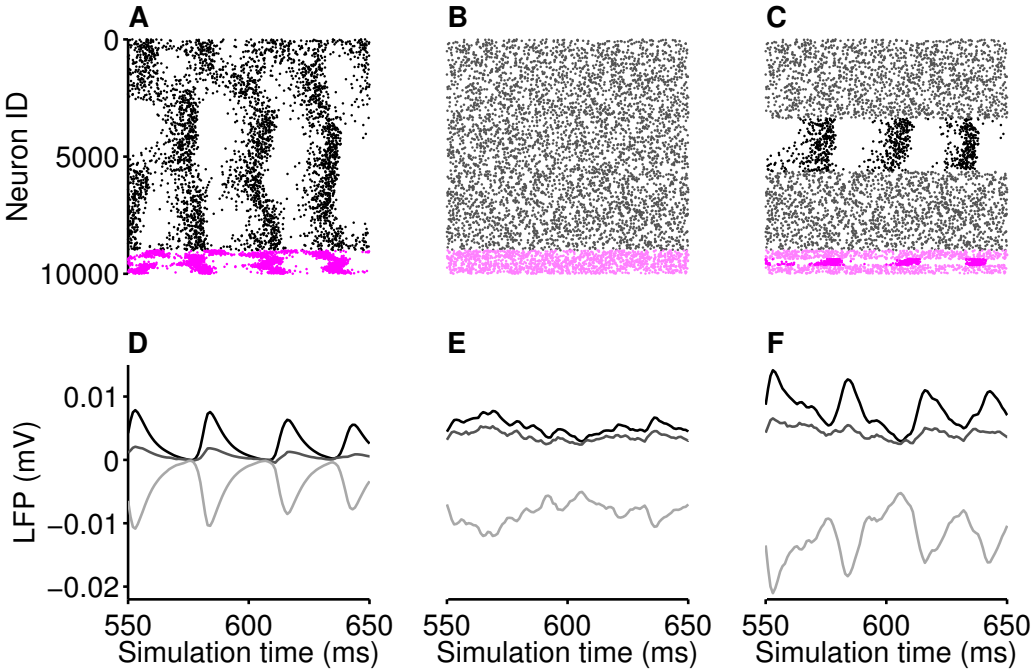

Figure 5.2: Spike rasters and example simulated LFPs. (A) Cell spiking during simulated gamma oscillation. (B) Cell spiking during random activity. (C) Combined gamma and random activity, to simulate spatially localised gamma generation. Pyramidal cell spikes shown in (faded) black, basket cell spikes in (faded) magenta. (D-F) Simulated LFP recorded at $x = 1100$ microns, $z = 500$ microns (black trace), $z = 200$ microns (medium grey trace) and $z =$ -50 microns (light grey trace). (D) LFP from central gamma generating region only. (E) LFP from surrounding random activity only. (F) Combined total LFP.

- Magnitude: standard deviation of the LFP recording.

- Power: variance of the LFP recording.

- PSD (power spectral density) or power spectrum: the power in the LFP as a function of frequency. We calculated the PSD using the Thomson multitaper method (Percival and Walden, 1993; Thomson, 1982), with a time-bandwidth product of 3, as implemented in the Matlab Signal Programming Toolbox. The integral of the PSD over the frequency range of the signal is proportional to the LFP power.

- Gamma power: LFP power in the gamma frequency band. This is found by integrating the PSD between 25 Hz and 40 Hz (though the gamma-band is usually defined up to 80 Hz, the peak frequency in our simulations was between 30 Hz and 35 Hz, so cutting-off at 40 Hz prevented harmonics of this fundamental frequency from being integrated as part of the gamma power).

- Gamma spatial range, or gamma visibility: after finding the PSD in a model $M\gamma$ with a gamma generating region (R_γ) surrounded by randomly firing neurons (R_R), we then find the PSD in a model M_R containing only randomly firing neurons (i.e. $X_\gamma = 0$ microns). If the lower 95% confidence interval of the PSD in M_γ overlaps the upper 95% confidence interval of the PSD in M_R between 25Hz and 40Hz at 3 frequency bins or more, we consider there to be a significant gamma component in the LFP (see Figure 5.3). The x-coordinate at which the oscillation first becomes insignificant, moving away laterally from the centre of the slice (at $x = 2000$ microns), is defined as the gamma spatial range, or gamma visibility. This is calculated for each z-depth of the electrode grid. We specify 3 or more frequency bins as the minimum required overlap for the following reason. We are comparing 15 frequency bins (between 25 Hz and 40 Hz). As the comparison is between 95% confidence intervals, we would expect them to overlap by chance on 5% of the bin comparisons. Considering a binomial distribution (they either overlap or they don't), if we make 15 bin comparisons then the probability of one chance overlap is 0.54, of two chance overlaps is 0.17 and of three chance overlaps is 0.036. The highest accepted p-value for the presence of a gamma oscillation in the signal, requiring 3 bins

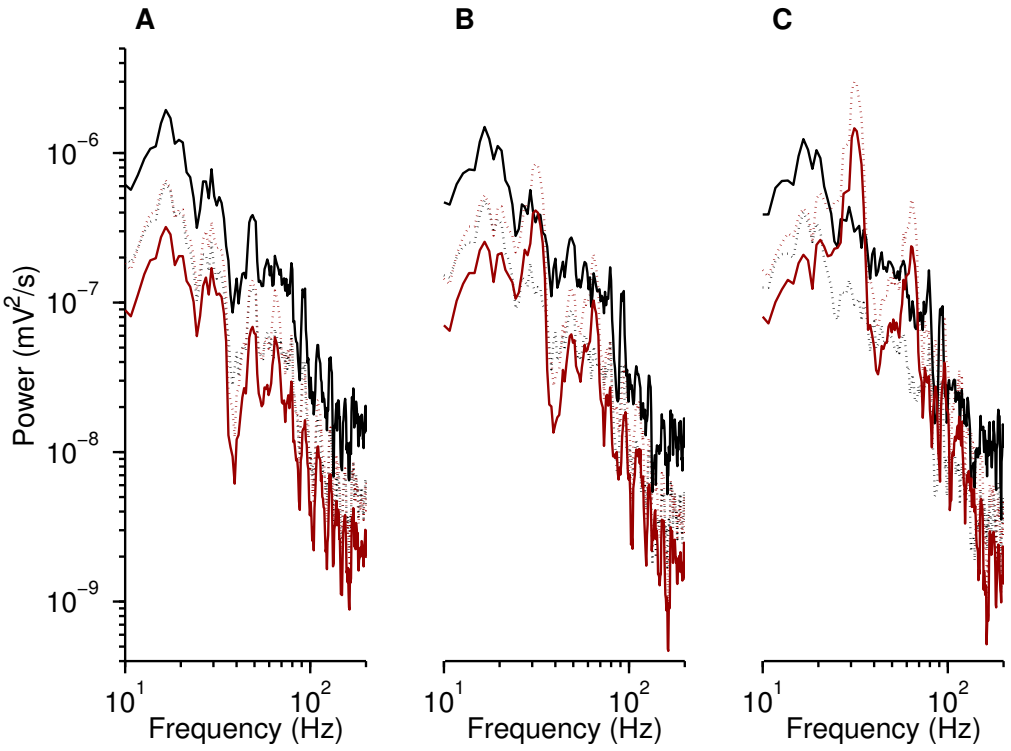

Figure 5.3: Method for measuring gamma oscillation LFP spatial range (shown here for LFPs measured at $z = 50$ microns). Solid red lines show the lower 95% confidence interval bound of the LFP PSD for the gamma model with $X_\gamma = 500$ microns; solid black lines show the upper 95% confidence interval bound of the LFP PSD for the entirely randomly spiking model (i.e. $X_\gamma = 0$ microns); dotted red and black lines show the LFP PSD estimate for $X_\gamma = 500$ microns and $X_\gamma = 0$ microns, respectively. (A) PSD for LFP measured at $x = 1300$ microns: no significant gamma. (B) PSD for LFP measured at $x = 1400$ microns: overlap of 2 frequency bins of the CI bounds - gamma not deemed to be visible. (C) PSD for LFP measured at $x = 1500$ microns: clearly significant gamma oscillation visible at this range.

to overlap, is therefore 0.036.

- Gamma coherence: the coherence in the 25 Hz to 40 Hz frequency band, measured between the LFP recorded at the centre of the slice (at x=2000 microns) and the LFP recorded at some distance along the x-axis away from the centre, at the same depth. The coherence spectrum between two LFP signals is calculated using the Matlab Signal Processing Toolbox's mscohere() function, and the gamma coherence is found by integrating the coherence between 25 Hz and 40 Hz.

- Power ratio, pow(Pre,Post↔All,All): the ratio between the power of the LFP due to the specified pre- and postsynaptic populations and the total LFP from all neurons. For example, the presynaptic pyramidal power ratio,

pow(P,A↔A,A), is equal to the power of the LFP resulting only from pyramidal cell firing divided by the power of the total LFP. The postsynaptic pyramidal power ratio, pow(A,P↔A,A), is the power of the total LFP divided by LFP resulting only from synaptic activity onto pyramidal neurons. A value of 1 means that the LFP generated by the specified pre- and postsynaptic groups contains the same power as the total LFP. It is possible for the power ratio to be >1, if other LFP contributions are out of phase with the pre- and postsynaptic populations under consideration.

- Correlation coefficient, CC(Pre,Post↔All,All): the correlation coefficient between the LFP signal due to the specified pre- and postsynaptic populations and the total LFP signal. CC(Pre,Post↔All,All) ranges from -1 (totally anti-correlated, i.e. the presynaptic neuron synapses onto the postsynaptic population actually reduce the total LFP power) to 1 (totally correlated).

- Phase-inversion: if the depth cross-correlation between the LFP and the LFP at the bottom of the slice (i.e. as compared to the electrode at the same x-location and at $z = 0$ microns) is negative, then it is considered to be phase inverted (see Figure 5.4).

- Phase-inversion depth: the depth at which a phase inversion is first observed, if moving upwards along an electrode column from the bottom of the slice.

5.3 Results

5.3.1 Population contributions to the LFP

Before further investigating the total LFP across space in the layer 2/3 models, we first wanted to establish the contribution of each neuron group to the total LFP. To enable us to look at each population's presynaptic LFP contribution, we used the spike-import feature of VERTEX to load and replay spikes fired by each group in turn, recording the resulting LFP resulting from each group's activity. We further subdivided the LFP signal by postsynaptic contribution, by setting all synapse strengths to zero, except for those onto the postsynaptic population of interest.

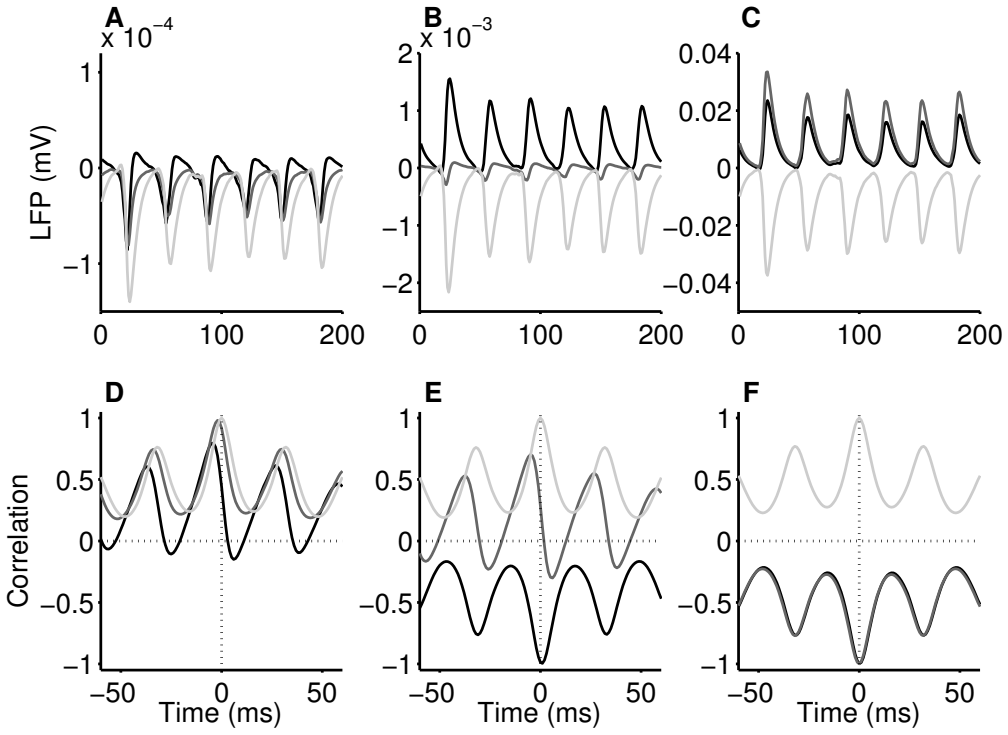

Figure 5.4: LFPs measured during gamma oscillation in 500 micron wide central network section, no additional random spiking. Each subfigure shows results at depths of $z = 500$ microns (dark grey), $z = 250$ microns (medium grey) and $z =$ -50 microns (light grey). All measurements from the 100 micron deep slice. (A) LFP at slice edge, $x = 0$ microns. (B) LFP at $x = 1000$ microns. (C) LFP at centre of gamma generating region, $x = 2000$ microns. (D) Cross-correlation of signals recorded at each depth with signal at $z =$ -50 microns ($x = 0$ microns). No phase inversion. (E) As (D), but $x = 1000$ microns. Phase inversion between $z = 250$ microns and $z = 500$ microns. (F) As (D), but $x = 2000$ microns. Phase inversion between $z =$ -50 microns and $z = 250$ microns.

We first checked our assumption that synapses onto pyramidal neurons dominated the LFP by calculating pow(A,P↔A,A). For every slice depth and gamma-region width, mean pow(A,P↔A,A) over all electrode positions was 0.99 or more, with a maximum standard deviation of 0.03. The lowest ratios were at the phase-inversion depth (see below). We also measured CC(A,P↔A,A): the mean CC(A,P↔A,A) over all electrode positions was 1.00, with a maximum standard deviation of 0.01, again for every slice depth and gamma-region width. We therefore assume that, even if using incorrect interneuron model sizes underestimates their contribution to the LFP, this will have had a minimal impact on our calculations because of the magnitude of the pyramidal cell dominance.

We next investigated the LFP contribution from each presynaptic population. We show results from the model with layer depth 200 microns and $X_\gamma = 500$ microns, but the general patterns are similar for each model size. The top row of Figure 5.5 shows (A) pow(P,A↔A,A) (B) pow(B,A↔A,A) and (C) pow(NB,A↔A,A) over the space of the model, when only $R\gamma$ was spiking (i.e. the surrounding neurons were silent). The power ratio was calculated at each electrode, with bicubic interpolation used to smooth the resulting image (Matlab interp2() function). The LFP is clearly dominated by basket cell activity, across the whole space of the model, while pyramidal activity contributes a small amount in a small region above the soma-layer. This region is where the basket cell activity contribution is also slightly reduced. The bottom row of the figure shows (D) CC(P,A↔A,A) across the slice, (E) CC(B,A↔A,A) across the slice, and (F) CC(NB,A↔A,A) across the slice. Again, for each plot CC was calculated at each electrode and smoothed with bicubic interpolation. As expected from the power ratio results, CC(B,A↔A,A) is very close to 1 across the whole slice, except in the small regions above the phase-inversion depth, which correspond to where pow(B,A↔A,A) is also reduced. CC(P,A↔A,A) is mostly close to 0, except in the locations in which pow(P,A↔A,A) is slightly elevated. Finally, CC(NB,A↔A,A) is negative over the whole space of the model, indicating that while the NB interneurons spike roughly synchronously with the B interneurons, their LFP contribution is out of phase with the total LFP. CC(NB,A↔A,A) increases in the same region that CC(P,A↔A,A) increases, but remains negative.

Next, we examined the LFP contributions in the model including randomly spiking neurons in the regions surrounding R_γ. We only present results for surround-

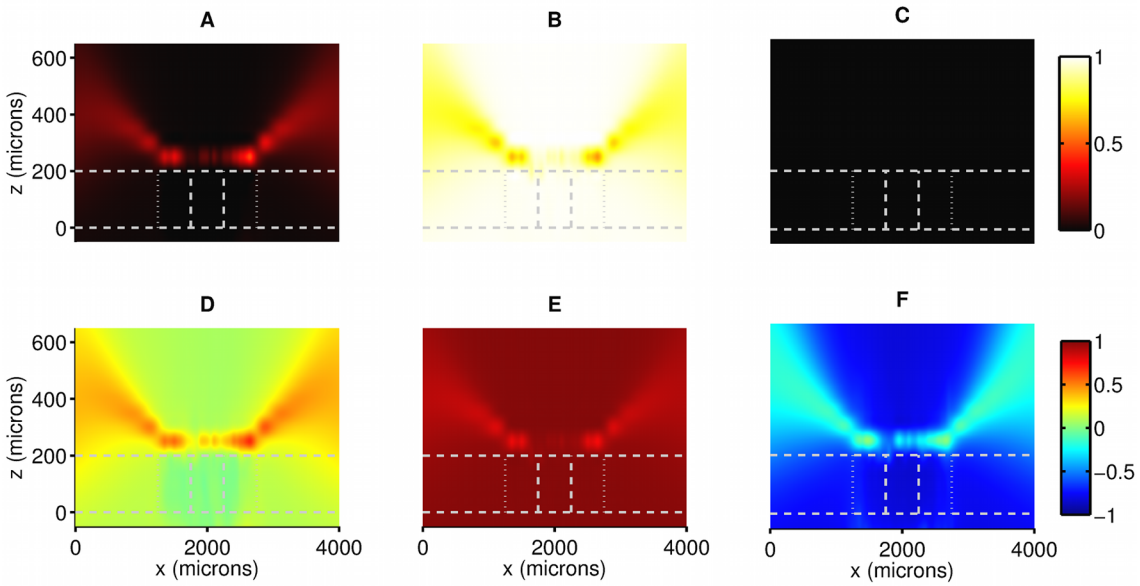

Figure 5.5: Presynaptic contributions to LFP for each group during a gamma oscillation in the centre of the slice, with the surrounding neurons not firing at all. Dashed grey lines show the slice borders and the extent of the gamma-generating network region; dotted grey lines show the maximum extent of the gamma-generating region's axonal arbours. (A) Ratio of the LFP power due to excitatory synaptic activity to the total LFP power, across the space of the model. (B) Ratio of the LFP power due to basket interneuron synaptic activity to the total LFP power, across the space of the model. (C) Ratio of the LFP power due to non-basket interneuron synaptic activity to the total LFP power, across the space of the model. Basket interneurons dominate the LFP power, across the whole space of the slice. (D) Correlation coefficient between the LFP due to pyramidal synapses and the total LFP. (E) Correlation coefficient between the LFP due to basket interneuron synapses and the total LFP. (F) Correlation coefficient between the LFP due to non-basket interneuron synapses and the total LFP. The basket interneuron contribution is highly correlated across the slice, while non-basket interneurons show a negative correlation, meaning that their LFP contribution would interfere negatively with the gamma oscillation.

ing regions with high random activity (5 Hz firing rate per neuron); the results are similar with low random activity (1 Hz firing rate per neuron), but the influence of R_γ extends further to the left and right. Figure 5.6 shows the same plots as Figure 5.5, for a gamma-generating region surrounded by randomly firing neurons. The power ratio and CC maps both reveal a clear boundary between LFPs that are primarily influenced by the gamma oscillation, and LFPs primarily influenced by the random activity. This boundary occurs close to the axonal arborisation extent of $R\gamma$ in both the soma layer and at the apical dendrite level, but widens towards the top of the measurement area. Basket cell synaptic activity still dominates the LFP even outside of the gamma-generating region, though pow(B,A↔A,A) and CC(B,A↔A,A) are reduced outside the gamma-generating region, while pow(P,A↔A,A), CC(P,A↔A,A) and CC(NB,A↔A,A) are increased. The reduction in pow(B,A↔A,A) and CC(B,A↔A,A), and corresponding increases in pow(P,A↔A,A) and CC(B,A↔A,A) are more pronounced within the soma layer than at the apical dendrite level. The variability across space of these values is also more pronounced in the soma layer - the values at the apical dendrite level appear to be spatially smoothed. Just above the soma-layer, pow(P,A↔A,A) and CC(P,A↔A,A) make the largest contribution to the LFP outside the gamma-generating region, while pow(B,A↔A,A) and CC(B,A↔A,A) are greatly reduced. pow(NB,A↔A,A) shows a very small increase outside the gamma-generating region, while CC(NB,A↔A,A) becomes positive, though close to zero.

These results show clearly that different presynaptic neuron groups will contribute in different ways depending on the network state - hardly surprising given the range of firing rates and correlation strengths between these different groups. It is interesting to note, though, how different groups, though firing synchronously, can interfere positively or negatively with the overall LFP signal depending on the placement of their synapses on the pyramidal cell dendrites. More surprising is the dominance of basket interneuron synapses in the LFP signal during random activity. The LFP contribution from spiking interneurons has only recently been investigated experimentally (Bazelot et al., 2010; Menendez de la Prida and Trevelyan, 2011; Oren et al., 2010; Trevelyan, 2009; Glickfeld et al., 2009). Our results regarding the dominance of basket interneuron firing on the LFP in R_γ fit with the experimental results reported by Trevelyan (2009), who found a high correlation between LFPs

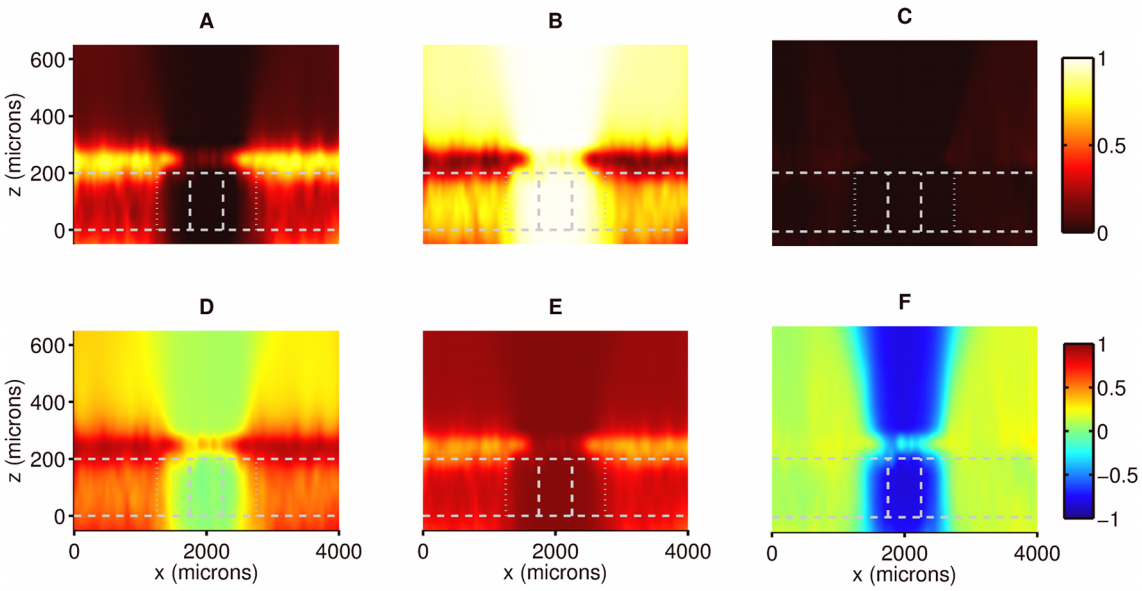

Figure 5.6: Presynaptic contributions to LFP for each group during a gamma oscillation in the centre of the slice, with the surrounding neurons firing randomly at 5Hz. Dashed grey lines show the slice borders and the extent of the gamma-generating network region; dotted grey lines show the maximum extent of the gamma-generating region's axonal arbours. (A) Ratio of the LFP power due to excitatory synaptic activity to the total LFP power, across the space of the model. (B) Ratio of the LFP power due to basket interneuron synaptic activity to the total LFP power, across the space of the model. (C) Ratio of the LFP power due to non-basket interneuron synaptic activity to the total LFP power, across the space of the model. Basket interneurons dominate the LFP power, most prominently in the gamma-generating region but also in the randomly firing region. Their contribution is lowest at the point of phase-inversion, indicating that the phase inversion is primarily due to currents from basket interneuron synapses. (D) Correlation coefficient between the LFP due to pyramidal synapses and the total LFP. (E) Correlation coefficient between the LFP due to basket interneuron synapses and the total LFP. (F) Correlation coefficient between the LFP due to non-basket interneuron synapses and the total LFP. In the randomly firing regions, all contributions are positively correlated with the total LFP, even if to a small degree in the case of non-basket interneurons. The non-basket interneurons are negatively correlated with the total LFP in the gamma-generating region, as in the previous figure. The LFP contribution of different neuron types clearly varies between different network states.

and simultaneously recorded inhibitory postsynaptic potentials in pyramidal neurons, and Oren et al. (2010), who showed that inhibitory currents in the perisomatic region of pyramidal neurons were the LFP generator during hippocampal gamma oscillations. The large contribution of basket cells to the LFP during random firing was not expected, but is consistent with the results reported by Bazelot et al. (2010), who found in rat hippocampal slices that single interneuron spikes evoked measurable LFPs, but single pyramidal neuron spikes did not.

Menendez de la Prida and Trevelyan (2011) propose some explanations regarding the contribution of inhibitory currents to the LFP, suggesting that high visibility may arise from (1) more synchronous firing patterns due to gap junction coupling, (2) their higher synaptic release probabilities compared with glutamatergic synapses, (3) the location of glutamatergic synapses on dendritic spines, which create smaller currents than synapses located directly on the dendrites or soma, (4) the density of connections each basket cell makes, and (5) the localisation of basket cell synapses to the perisomatic region of the pyramidal cells. We can estimate the impact of these points by considering the LFP generated by a theoretical neuron firing a Poisson spike train at a set rate. We assume that this neuron makes a large number of connections and its connectivity pattern is radially symmetrical, so we can neglect the influence of spatial variation in connectivity. We also assume that it makes all its synapses at the same location on all its targets, and that its targets are identical, so we can ignore the influence of synapse location and neuron morphology. Starting with point 4, if we increase the number of synapses (i.e. the synaptic density), maintaining radial symmetry in the connectivity, then the LFP magnitude will increase proportionally to the number of synapses. In this simplified situation, this is also equivalent to increasing or decreasing the synaptic weight (points 2 and 3); the LFP magnitude changes proportionally to the weight. The LFP magnitude also increases proportionally to the square root of the synaptic decay time constant. Instead of (or in addition to) increasing the number of synapses, we could increase the number of neurons by introducing extra identical neurons at the same location. If each fires an independent Poisson spike train, then the total LFP magnitude will increase proportionally to the square root of the number of neurons, i.e. the number of Poisson processes (Benedek and Villars, 2000). If each neuron fires at the same time (point 1), then the effect of adding more neurons will have the same effect as

increasing the synaptic weights, so the LFP magnitude will increase proportionally to the number of neurons. So, our estimate for the magnitude of the LFP $\left|\tilde{\phi}\right|$ due to a presynaptic population of N theoretical neurons is

$$\left|\tilde{\phi}\right| \propto \begin{cases} n_{syn} \cdot w_{syn} \cdot \sqrt{\tau_{syn}} \cdot \sqrt{N} & \text{if firing is uncorrelated} \\ n_{syn} \cdot w_{syn} \cdot \sqrt{\tau_{syn}} \cdot N & \text{if firing synchronously} \end{cases}, \tag{5.1}$$

where n_{syn} is the number of synapses per theoretical neuron, w_{syn} is the synaptic weight, and τ_{syn} is the synaptic decay time constant. The one factor mentioned above that is excluded from this model is the location of the synapses on the postsynaptic neurons' dendrites. Note that the LFP magnitude or power estimated in this way is a different measure than the proportional power plotted in Figures 5.5 and 5.6. The proportional power takes into account the contributions of other populations to the total LFP, while the rough LFP magnitude estimate considers an individual presynaptic population.

For our model, P cells and NB cells can make synapses onto the same compartments of their synaptic targets, so synapse location should not affect the relative LFP magnitude from each of these populations. Considering only synapses onto pyramidal neurons (as these dominate the LFP), the P cell LFP magnitude in the uncorrelated firing regions should be proportional to $694 \times 1 \times \sqrt{2} \times \sqrt{9000} = 9.3 \times 10^4$, while the NB cell LFP magnitude should be proportional to $549 \times 1.8 \times \sqrt{6} \times \sqrt{1000} = 7.7 \times 10^4$. In other words, the LFP magnitude due to P cell spiking will be roughly 1.2 times the LFP magnitude (1.5 times the power) due to NB cell spiking. Discounting the different localisation of B cell synapses, the LFP magnitude due to their firing should be proportional to $872 \times 3.5 \times \sqrt{6} \times \sqrt{1000} = 23.6 \times 10^4$, or 2.5 times the P cell LFP magnitude. However, if our synaptic weight assumptions are wrong, and B cell & P cell synapses are in fact similar in strength, then we would predict that, without including synaptic localisation, the B cell LFP magnitude would be 0.73 times that of the P cell LFP magnitude (0.53 times the power).

We investigated this possibility by scaling the LFP contributions from each pair of populations by relevant synaptic weights - equivalent to scaling all the weights to be equal to 1 due to the linearity of the model. The results on the LFP contribution from each group in this case are shown in figure Figure 5.7.

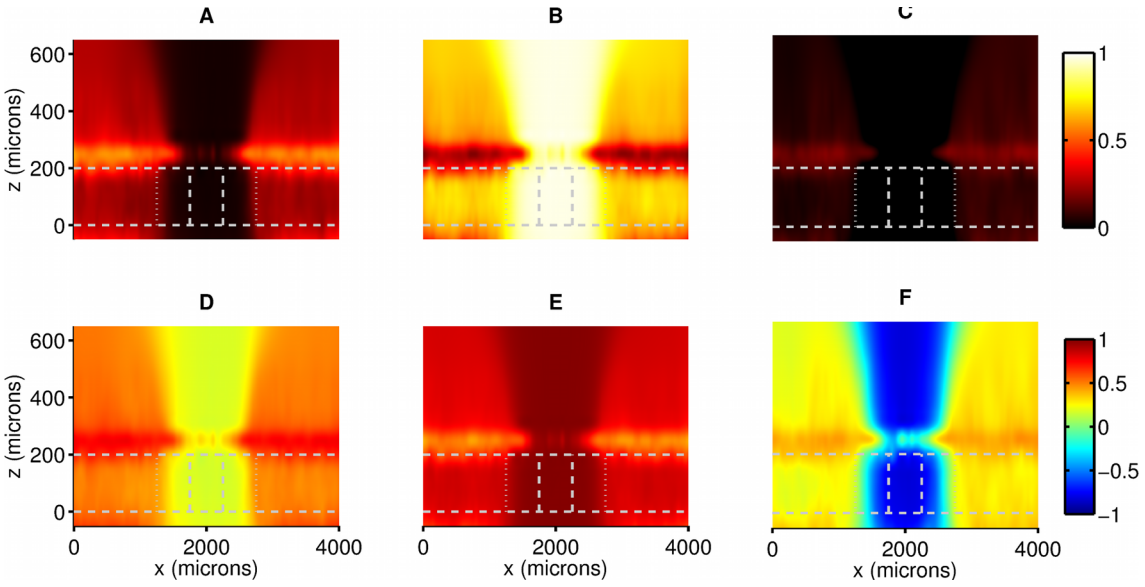

Figure 5.7: Presynaptic contributions to LFP for each group during a gamma oscillation in the centre of the slice, with the surrounding neurons firing randomly at 5Hz, as in the previous figure. For these results, all synaptic strengths were scaled to be equal, so that we could see the LFP contribution of each presynaptic neuron group independently of synapse strength parameters. Results are similar to the previous figure, except that non-basket interneurons contribute more to the LFP in the randomly spiking region. The basket interneuron contribution above the phase inversion level is also reduced compared to the previous figure; this is because basket interneuron synapses were 3.5 times stronger than pyramidal neuron synapses in that figure. Interestingly, though, basket interneurons remain the primary LFP contributor, suggesting that our previous results were not a result of our incorrect synaptic strength parameters.

The pattern of power ratios and CCs shown in Figure 5.7 is similar to that in Figure 5.6 with LFPs calculated using the original synaptic weights. The most obvious differences are a slightly increased pow(NB,A↔A,A) and CC(NB,A↔A,A) outside of the gamma dominated region and a slightly reduced pow(B,A↔A,A) and CC(B,A↔A,A), with corresponding increases in pow(P,A↔A,A) and CC(P,A↔A,A), just above the soma-layer outside of the gamma dominated region. The variability in the power ratio in the soma layer is also reduced.

Our results demonstrate that the location of the basket cell synapses at the perisomatic region of their pyramidal targets allows the activity of basket interneurons to dominate the LFP during random firing. This is true even when the LFP is measured at the level of the pyramidal neuron apical dendrites, which are primarily targeted by other pyramidal neurons and non-basket interneurons. The level of basket cell dominance increases during gamma oscillations to the extent that contributions from the firing of other neurons has a negligible influence on the LFP signal. This is because the basket cells fire with the greatest synchrony during the oscillation. Though we did not consider variable firing rate in our simplified analysis above, the rate of basket interneuron firing can be an order of magnitude greater than pyramidal neuron firing during gamma oscillations, which will also increase their LFP dominance.

5.3.2 Phase inversion

Figure 5.8 shows raw LFP traces at the central electrode column (i.e. for all electrodes at x=2000 microns), along with gamma power and phase as functions of depth, for the model with a soma-layer depth of 300 microns and gamma-generating region width of 1000 microns, showing a phase inversion at 300 microns above the bottom of the soma layer. Two peaks in the gamma power are apparent either side of the phase inversion depth: one at 100 microns and a second at 450 microns. We will return to the power in the next section.

Our gamma oscillation model provides theoretical support for the observation made in (Cunningham et al., 2003) (see Figure 5.9): the spiking activity in the model matches that observed experimentally, and the inhibitory synaptic currents are localised to the perisomatic region of the pyramidal neurons, causing the return current to originate primarily from the apical dendrite. Also as observed experi-

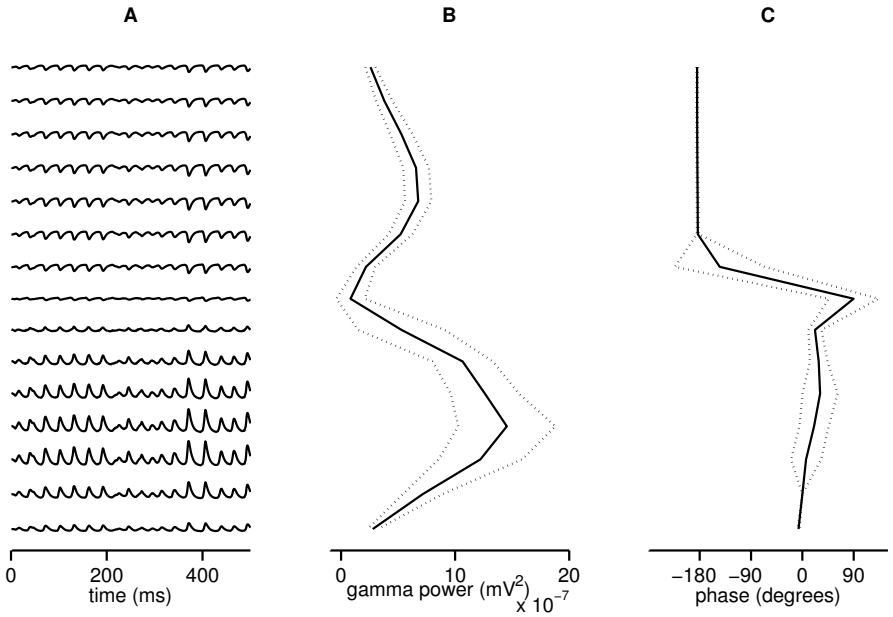

Figure 5.8: Raw LFP traces from a single run, mean power and mean phase over nine runs, measured at the central electrode column. $X_\gamma = 1000$ microns, soma layer depth = 300 microns. A: raw LFP traces from top electrode (z=650microns) to bottom electrode (z=-50microns). B: Gamma power at each electrode ($\pm$ 2 SEM). C: phase relative to the electrode at the bottom of the soma layer ($\pm$ 2 SEM).

mentally (Dickson et al., 2000, - see below), the power of the gamma oscillation is greatly reduced at the point of phase inversion, as the influence on the LFP from the current sinks and sources approximately balances.

We can make a further observation from our model regarding the location of the phase inversion point at a particular x-location, Φ_x^-, when we measure the LFP at a greater lateral distance from the gamma generating area: Φ^- occurs higher up the z-axis the further we measure laterally from the gamma generating region (though note that this only occurs in the model with no random surround firing; as such, this observation would be difficult to test experimentally). To create a profile of the phase inversion location measured across the slice for each size of gamma-generating region, we found the phase-inversion depth for each electrode column. The results of this analysis are shown in Figure 5.10, plotted against the basket interneuron power ratio pow(B,A$\leftrightarrow$A,A) to show how the reduction in basket cell influence matches the measured phase inversion depth.

The spatial profile of the phase inversion is similar for each condition: the phase inversion occurs higher further away from the gamma generating region. Φ_x^- reaches a minimum Φ_{min}^- outside of the boundaries of the gamma generating region for each

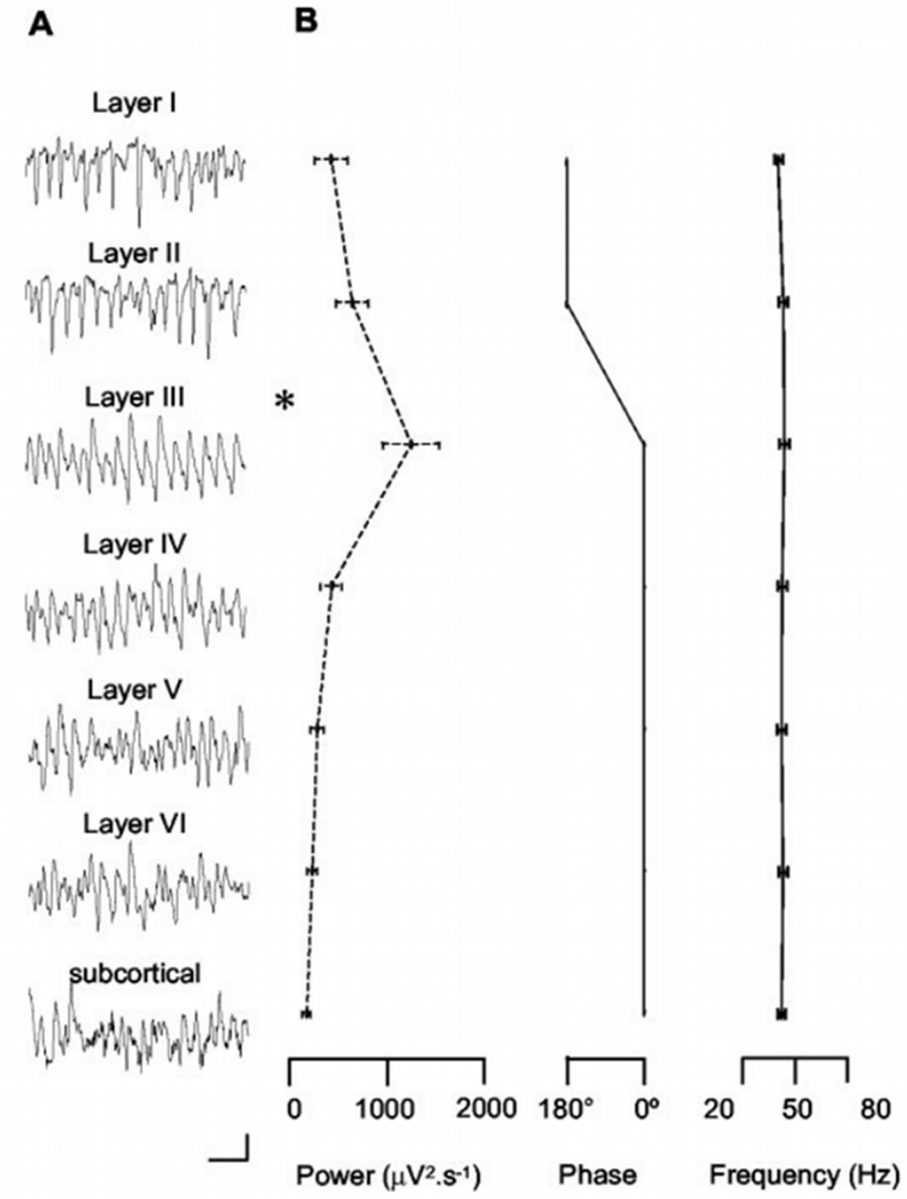

Figure 5.9: Figure 3 from (Cunningham et al., 2003), showing LFP recordings from slices of rat medial entorhinal cortex with gamma oscillations induced by bath kainate application. (A) shows the raw LFP traces measured in different cortical layers. (B) shows the power in the gamma frequency band (in this case, between 20Hz and 80Hz), phase, and peak frequency as a function of cortical depth. Scale bar in A shows 50 milliseconds horizontally and 50 microVolts vertically.

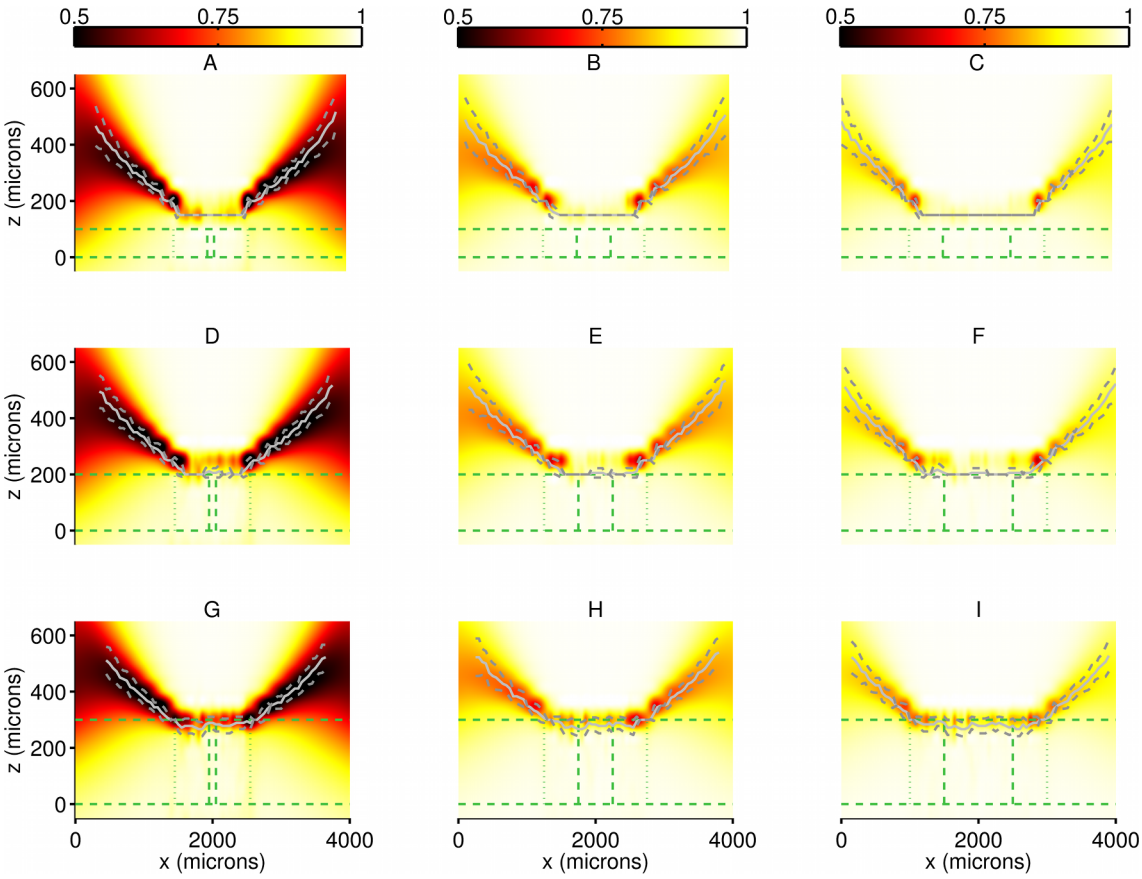

Figure 5.10: Phase inversion depth plotted for each soma-layer depth and $X\gamma$, over the relevant pow(B,A↔A,A) map. Phase inversion depth is the light line, with the standard deviation plotted as darker grey dashed lines. Dashed green lines show soma layer and R_γ boundaries, dotted green lines show maximum extent of synapses from R_γ. In many cases, the phase inversion depth plot does not extend to the model edges at $x = 0$ and $x = 4000$ microns. This indicates that a phase inversion was not detected at these electrode columns.

X_Γ, though this minimum is proportionally widest for smaller X_γ (for example, at the 100 micron layer depth, $\Phi_x^- = \Phi_{min}^-$ from $x = 1500$ microns to $x = 2500$ microns for $X_\Gamma = 100$ microns, while $\Phi_x^- = \Phi_{min}^-$ from $x = 1200$ microns to $x = 2800$ microns for $X_\Gamma = 1000$ microns i.e. the proportional width of the minimum is $(2500 - 1500)/100 = 10$ for $X_\gamma = 100$ microns and $(2800 - 1200)/1000 = 1.6$ for $X_\gamma = 1000$ microns).

The change in depth, and eventual disappearance of, the phase inversion with increasing distance from the gamma generating zone is a result of the distance dependence of the LFP and the particular spatial spread of the membrane currents along the pyramidal cells' dendrites. The largest membrane current in each pyramidal cell is at the soma, where it receives strong inhibition from the synchronous spiking of basket interneurons. Much of the return current flows through the apical dendrite, with each apical compartment contributing to a fraction of the return current. When the LFP is measured close to the pyramidal cell somas, the current from the somas dominates, and when it is measured close to the apical dendrites, the currents from the apical compartments dominate, though the soma currents are larger. If we take a point near the apical dendrites, then move it laterally away from $R\gamma$, the distance between the point and the apical compartments increases at a greater rate than the distance between the point and the somas, so the somatic currents have an increasing influence on the measured LFP, thus causing Φ_x^- to increase.

5.3.3 Gamma power and range

Next, we investigated the spatial profile of the gamma power to see how it decayed with distance from the gamma generating region, and at what range a significant oscillation would be observed in the LFP given random surrounding activity. To investigate gamma range, we defined what we considered to be a significant oscillation, as specified in section 5.2.6. This definition of LFP range is different from the definition used in Chapter 2 - and in (Lindén et al., 2011) - because of the different experimental setup: in our current experiments we wanted to measure how far outside the generating network a particular pattern of activity could be detected in the LFP, while in Chapter 2 we were measuring the maximum distance from a recording point that neurons can influence the LFP magnitude.

To visualise how the depth of the soma layer and width of the gamma-generating

region affect the spatial profile of the gamma power, we plot gamma power maps over the space of the model. Figure 5.11 shows the variation in gamma power over space for each model condition, as well as the range of the oscillation away from the gamma-generating region. The gamma power is always greatest in the soma layer in the gamma-generating region, where the density of synapses receiving spikes from neurons entrained to the gamma oscillation is greatest. The power in this region also decreases with layer thickness: the number of neurons in the layer is constant, so the effect of increasing the layer thickness is to reduce the neuronal density. Gamma power decreases with decreasing density because the neurons are more spread out, meaning that the average distance of the pyramidal cell somas to any given recording point is larger. However, the power of the secondary gamma peak above the phase-inversion depth does not decrease in proportion to the decrease at the soma level, meaning that the difference in power between the two peaks decreases with increasing soma-layer thickness.

This effect is due to how the current sources and sinks are localised on the pyramidal cells. To illustrate this, consider two P cells in the model receiving identical inhibitory synaptic input at their somas, with their somas at points (x_1, z_1) and (x_2, z_2) and two electrode locations half way between these neurons at different cortical depths - one at the soma level and one at the apical dendrite level - as illustrated in Figure 5.12 A. The dominant LFP contribution at the bottom electrode will be from the soma compartments, so will be proportional to $1/d_1 + 1/d_2$, while the dominant LFP contribution at the top electrode will be from the apical dendrites and proportional to $1/d_3 + 1/d_4 + 1/d_5 + 1/d_6$. As we increase the soma depth as illustrated in Figure 5.12 B, we increase the distance $z_2 - z_1$, so d_2 will increase proportionally to $\arcsin(z_2 - z_1)$. However, while d_6 increases as $z_2 - z_1$ increases, d_5 *decreases*, so that the LFP change at the top electrode will be much smaller. In the model with a X_γ= 1000 microns, for example, the peak gamma power at the dendrite level is only 1.4 times as great when the soma-layer depth is 100 microns as it is when the soma-layer depth is 300 microns, while at the soma level it is 3.4 times as great. We list the ratio of peak gamma power in the dendritic region to peak gamma power in the soma region for each model configuration in Table 5.3, which shows that this effect is consistent for each width of gamma-generating region.

The results in Table 5.3 suggest a possible mechanism by which the gamma

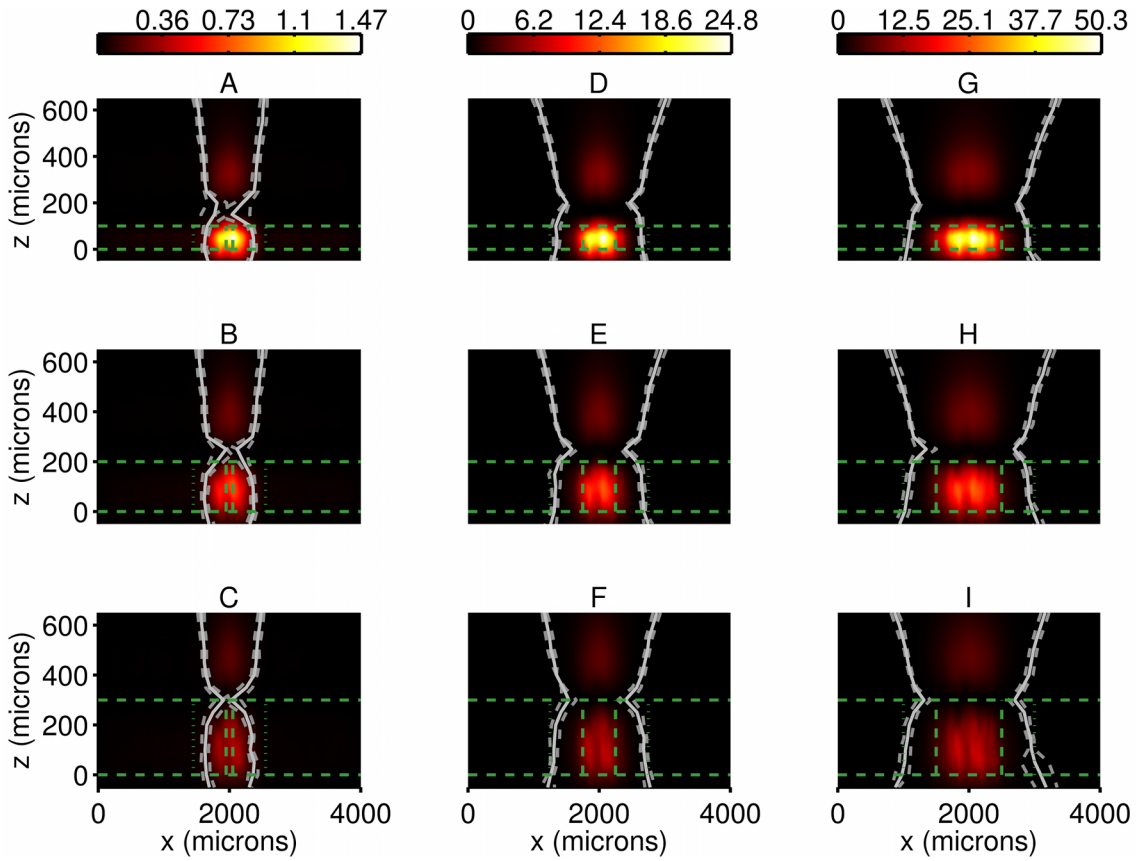

Figure 5.11: Magnitude and spatial range of the gamma oscillation for various values of X_γ and layer depths. Far left: $X_\gamma = 100$ microns, centre: $X_\gamma = 500$ microns, far right: $X_\gamma = 1000$ microns. Top row: thickness of soma layer = 100 microns, centre row: thickness of soma layer = 200 microns, bottom row: thickness of soma layer = 300 microns. Note different scale bars. Gamma range is plotted in light grey solid lines, with standard deviation indicated by grey dashed lines. Dashed green lines show soma layer and R_γ boundaries, dotted green lines show maximum extent of synapses from R_γ. Note the error bars showing a large variability in range between the different network instances for the thin gamma generating region.

	100 microns	**500 microns**	**1000 microns**
100 microns	0.18	0.19	0.19
200 microns	0.30	0.32	0.32
300 microns	0.47	0.47	0.47

Table 5.3: The ratio of maximum mean gamma power above the soma-layer to maximum mean gamma power in the soma-layer, for each model configuration.

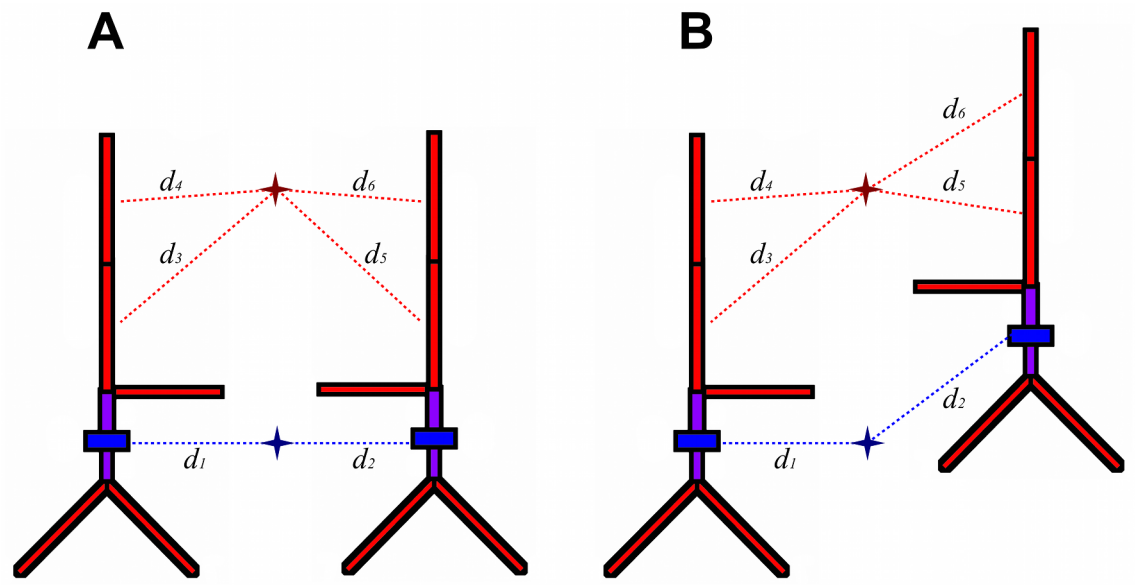

Figure 5.12: Simplified explanation of why the change in gamma power at the apical dendrite level is small compared to the change at the soma level for increasing soma-layer depth. Colours represent the direction of current flow: blue for outward currents (from inhibitory synapses) and red for inward currents (the return current). (A) Model with a small soma layer depth. The relative distances between equivalent compartments on the two neurons to the two electrode points are the same. (B) Model with increased soma layer depth. The LFP at the bottom electrode will be reduced as the d_2 increases. However, while d_6 increases, d_5 decreases, so the overall effect on the LFP magnitude at the top electrode is small compared with that at the bottom electrode.

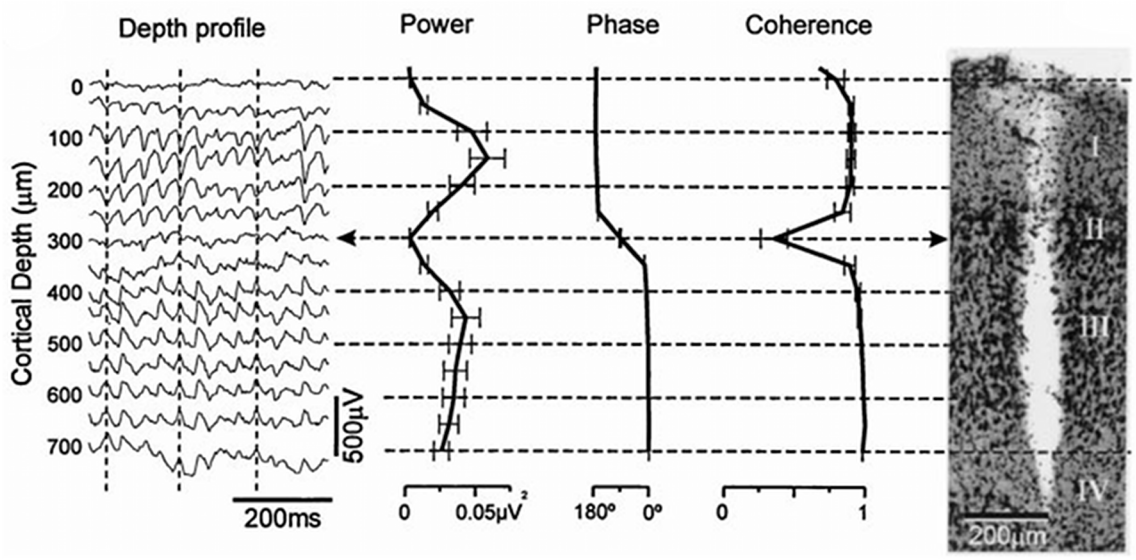

Figure 5.13: Figure 1 C from (Dickson et al., 2000) showing LFP recordings from a 16 site linear silicon probe placed in medial entorhinal cortex (mEC) in the guinea pig whole brain preparation. Gamma oscillations were evoked with arterial perfusion of carbachol. Raw LFP traces are shown on the left, with power, phase and coherence measurements along the depth of the probe show in the middle, and the probe tract shown on the right.

oscillation power observed in the guinea pig whole-brain preparation by Dickson et al. (2000) was actually greater at the apical dendrite level - in layer 1 - than at the soma level (layer 2/3, see Figure 5.13). If the soma-layer depth in the medial entorhinal cortex where they measured was thicker than 300 microns, as it appears to be in Figure 5.13, then the reduction in gamma power at the soma level could possibly be large enough that the power at the apical dendrite level will be greater than at the soma level. Unfortunately we made this observation rather late in the day so did not have time to test our hypothesis in a model with a thicker soma layer, but this would be very easy to do as a next step in our research.

Figure 5.11 also shows the gamma oscillation range away from the centre of the slice, defined as the distance away from the central electrode column at which significant 25 Hz - 40 Hz power in the signal is no longer visible. The range is smallest just above the soma-layer, and increases above and below this depth. The standard deviation of the gamma spatial range is very small above the soma-layer, but large within it and very large at the phase-inversion depth. This was particularly the case when the gamma-generating region was only 100 microns wide. This width only contains 225 P cells, 25 B cells and 25 NB cells, so variation between model runs due to randomness in the connectivity, neuron placement, and spike timing

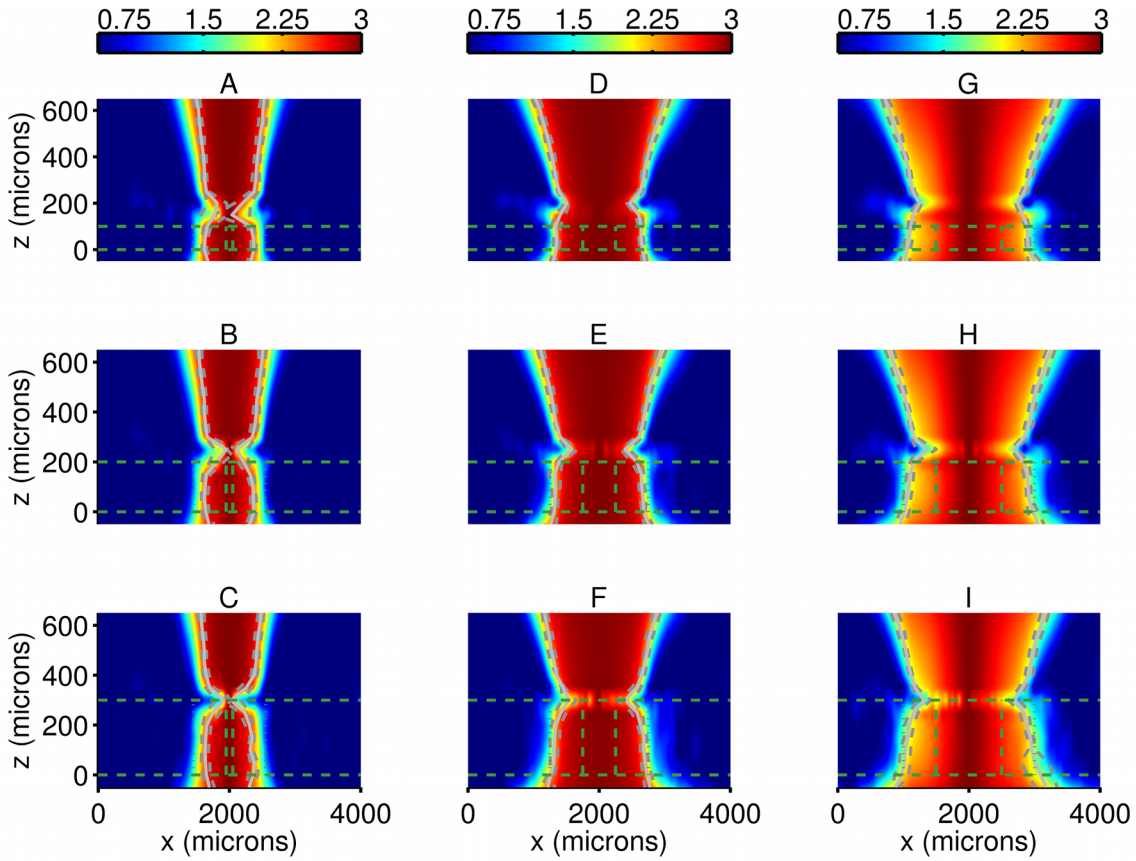

Figure 5.14: Gamma coherence (coherence in the gamma frequency band, measured in relation to the centre of the model) for various values of X_γ and layer depths. Far left: $X_\gamma = 100$ microns, centre: $X_\gamma = 500$ microns, far right: $X_\gamma = 1000$ microns. Top row: thickness of soma layer = 100 microns, centre row: thickness of soma layer = 200 microns, bottom row: thickness of soma layer = 300 microns. Gamma range is plotted in light grey solid lines, with standard deviation indicated by grey dashed lines. Dashed green lines show soma layer and R_γ boundaries, dotted green lines show maximum extent of synapses from R_γ.

will be much greater than in the models with larger X_γ. If the surrounding regions produced by chance an LFP with particularly low or high gamma frequency content, this would affect the gamma range to a greater extent than in models with larger gamma-generating regions. The reduced variance above the soma-layer indicates a spatial smoothing effect of the return current through the apical dendrites, given that the gamma oscillation is primarily created by B cell inhibitory currents at the soma level.

To visualise how rapidly the gamma oscillation disappears away from R_γ, we plotted the mean gamma coherence over the space of the slice for each model configuration, along with the gamma range values, in Figure 5.14 . The coherence profiles fit very closely with the gamma range profiles, decreasing to their minimum values

close to the gamma range. Interestingly, the coherence profiles for the 100 micron wide gamma-generating region do not exhibit the same degree of variability as the gamma range profiles. They show a rapid decrease down to the lowest value close to the edge of where the gamma power remains visible.

As an alternative visualisation of the coherence to get a better idea of the rate of decay away from the centre of the slice, we plot the coherence level across space for individual rows of electrodes (top, bottom and the row at the phase-inversion depth) for the three gamma-generating widths in Figure 5.15. All plots in Figure 5.15 show results for a model with a 200 micron soma depth. To see how the decay in coherence away from the centre of the slice changed with the level of activity in the surrounding regions, we also plotted the same measurements made when neurons in the surrounding regions were firing at 1Hz (Figure 5.15 D-F) rather than the 5Hz rate used previously (Figure 5.15 A-C). The effect on the coherence is to widen the central region of high coherence, with the rate of decay outside this central region similarly steep. The gamma range in the low surround-activity regime was greater at all depths, but otherwise showed a similar spatial profile. The cause of the additional coherence peak at the centre of the slice in the model with the 1000 micron gamma-generating region is unknown, but may be because there is some spatial variation in the gamma oscillation in the original gamma generating model (see Figure 5.2 A). Only when the gamma-generating region incorporates enough neurons does the effect of this variation become apparent in the hybrid model.

These results suggest that, when a gamma oscillation is observed in the LFP - at least during persistent gamma *in vitro* - we can be reasonably certain that the network generating the oscillation is close to where the electrode is positioned rather than being further away but spread out through volume conduction. This is important to consider when measuring lateral network synchrony using LFP recordings from MEAs (Ainsworth, 2013).

5.3.4 Frequency scaling

The final LFP feature we looked at in the model was the power spectrum, to investigate the frequency scaling properties at different points in the model. Figure 5.16 shows log-log plots of the normalised LFP power spectra at 20 electrode locations in the model with soma-layer depth 200 microns and X_γ= 500 microns. The gamma

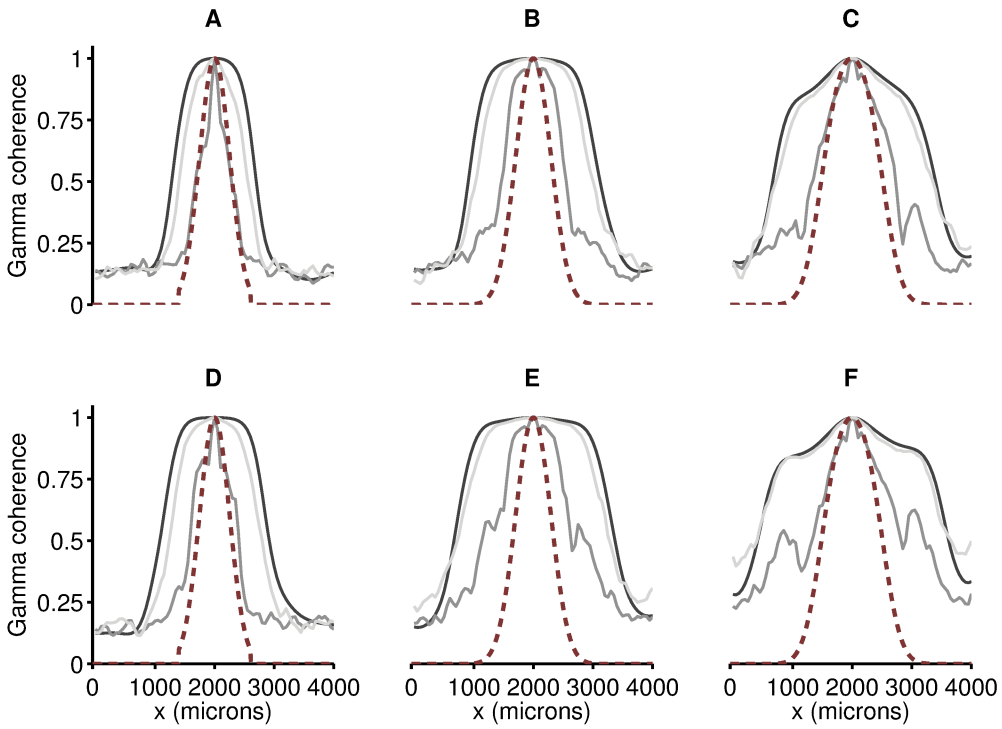

Figure 5.15: Gamma coherence in the model with soma-layer depth of 200 microns, measured at $z = 650$ microns (dark grey line), at $z = 250$ microns (medium grey line, at the phase-inversion depth), and at $z = -50$ microns (light grey line),. (A-C) Gamma coherence in the model with high surrounding random activity, for (A) $X_\gamma = 100$ microns, (B) $X_\gamma = 500$ microns, (C) $X_\gamma = 1000$ microns. (D-F) Gamma coherence in the model with low surrounding random activity, for (D) $X_\gamma = 100$ microns, (E) $X_\gamma = 500$ microns, (F) $X_\gamma = 1000$ microns.

peak can be seen clearly in the central column (x=2000 microns), as well as in the top row (z=650 microns) at x=1000 microns and x=3000 microns (as the gamma spatial range is greater at the top of the model; see above).

Several aspects of the LFP frequency scaling are apparent in Figure 5.16. Firstly, at all horizontal locations, the frequency scaling is shallower at the electrodes in the soma layer than at the electrodes above the soma layer. In the top row of electrodes, frequencies above ∼40 Hz show approximately $1/f^3$ (where f is the frequency) scaling except at the central column, while in the bottom row the frequency scaling is closer to $1/f^2$ except at the central column. The power spectra in central electrode column (in the gamma generating region) are dominated by the gamma oscillation and its harmonics up to ∼90 Hz. Above this, the frequency scales more steeply than in surrounding regions, steeper than $1/f^3$ in the soma layer and even steeper above.

The synaptic input to the pyramidal neurons that generates the majority of the LFP signal is primarily from basket interneurons, as described above. This means that it is localised to the perisomatic region, and the return current is low-pass filtered by the apical dendrites. This is clear from the different frequency-scaling exponents at the apical and soma levels in Figure 5.16. These results are consistent with those reported in (Lindén et al., 2010; Pettersen and Einevoll, 2008) regarding the origin of the low-pass filtering of LFPs. However, they contrast with the frequency scaling mechanism proposed by Bédard and Destexhe (2009), who suggest that the scaling effect is a result of neural tissue properties (see discussion in Chapter 2). To illustrate how synaptic input to pyramidal neurons would not produce the correct frequency scaling characteristics under *in vivo*-like, noisy input conditions, Bédard et al. (2010) simulated a "ball-and-stick" neuron (a neuron model with one soma compartment, and a chain of compartments extending vertically above the soma to represent the apical dendrite) receiving noisy synaptic input along the whole length of the dendrite, which they argue is representative of the conditions *in vivo*. The frequency scaling in the power spectrum of the simulated LFP showed no low-pass filtering, as the synaptic currents dominated over the passive return currents. However, as we have shown above, the localisation of basket interneuron synapses at the perisomatic region causes these inputs to dominate the LFP, even during random firing (which we suggest is also similar to the decorrelated state *in vivo*). When Bédard et al. (2010) simulated localised synaptic input to the ball-

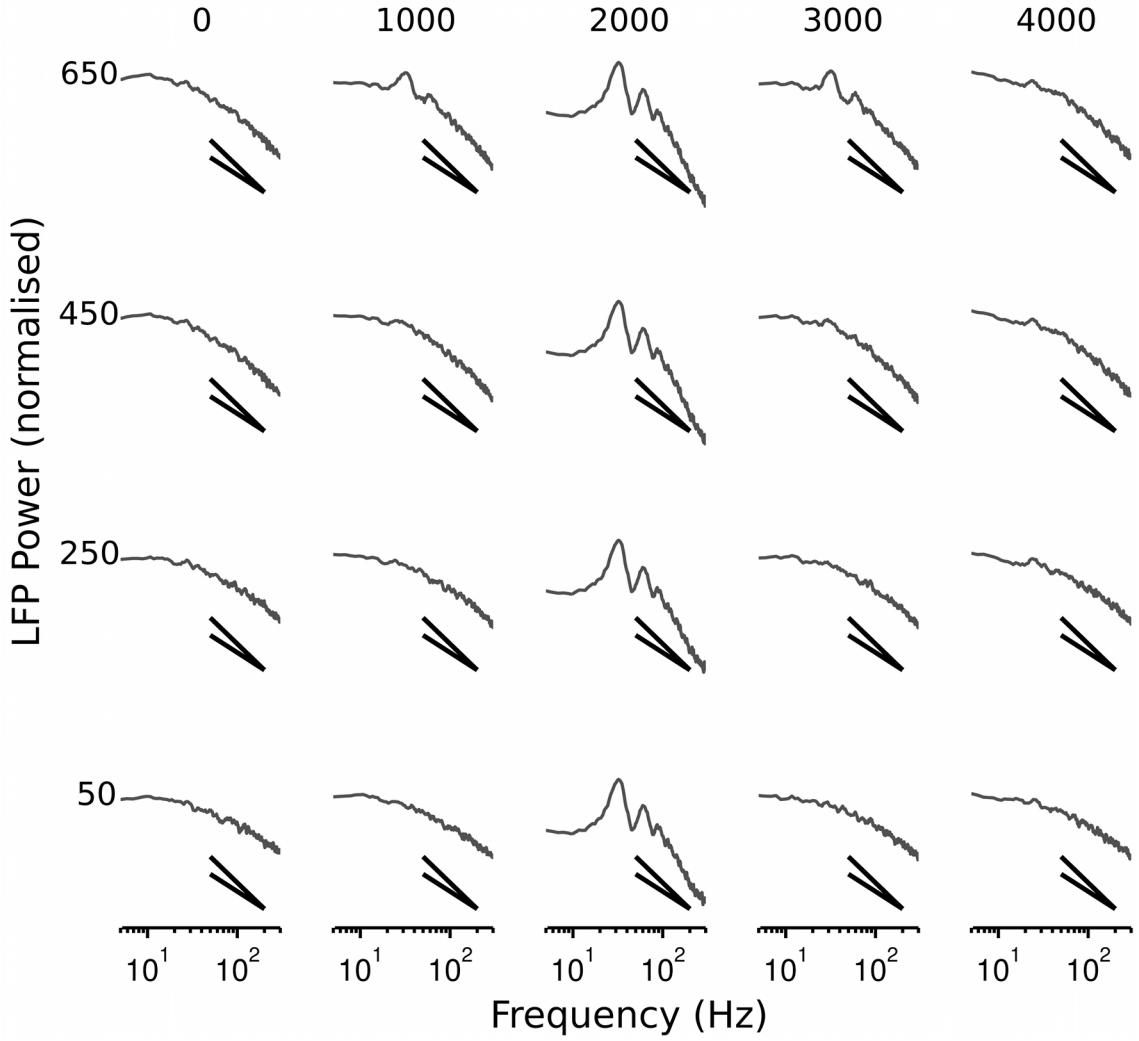

Figure 5.16: Mean power spectra of the LFP at 20 locations across the model (over all model runs) with gamma-generating region width 500 microns and soma layer depth 200 microns (log-log scale). Model coordinates for each electrode in the figure are given at the left and along the top of the figure, in microns. Scale bars show $1/f^2$ and $1/f^3$scaling. Note the shallower slope in the soma layer (bottom row) compared with the slope at the apical dendrite level (top two rows). The slope at higher frequencies also becomes steeper closer to the gamma-generating region.

and-stick model, strong low-pass filtering was observed, in agreement with (Lindén et al., 2010). We suggest that the perisomatic localisation of basket interneuron synapses would create such localised input even in decorrelated states *in vivo*.

The results reported in (Bédard and Destexhe, 2009) include comparisons between their proposed LFP model and example data recorded from cat *in vivo*, but do not report the electrode depth, so it is difficult to draw any firm conclusions from these observations. We make a comparison of our model results with the macaque slice data reported in Chapter 4, as we have the advantage of knowing where the electrodes were positioned. Figure 5.17 shows the LFP power spectra from electrodes 41 (in layer 1) and 42 (in layer 2/3) from the Utah array recording. As in our simulation, the frequency up to ∼90 Hz is dominated by the gamma oscillation and its harmonics. Above this frequency, the LFP power recorded at electrode 41 scales more steeply than $1/f^3$, while the LFP frequency scaling at electrode 42 is shallower, both than at electrode 41 and at the soma layer in our model. We suggest that this relatively large high-frequency content comes from the low-frequency contributions from spikes. Many neurons fire synchronously and rapidly during persistent gamma oscillations (particularly the chattering pyramidal neurons), so the LFP contribution from spiking activity is likely to be greater than during network states with low spike rates. Frequencies from extracellular spikes have been shown to contaminate the LFP at frequencies as low as 90 Hz (Schomburg et al., 2012; Zanos et al., 2011), and perhaps even lower in some cases (Waldert et al., 2013; Reimann et al., 2013). That the frequency scaling at the apical dendrite level (electrode 41) is still steep above 90 Hz suggests that the source of the increased high frequency component in the LFP at electrode 42 is localised close to electrode 42, and is filtered by the pyramidal apical dendrites. This fits with the hypothesis that this frequency content is due to spiking. Finally, the simulated LFP at the soma layer in our model, containing no spiking components, exhibits steeper frequency scaling than the data from electrode 42 right up to 300 Hz, but otherwise matches it.

Further work is required to investigate the causes of frequency scaling in the LFP during gamma oscillations - indeed during any activity regime. We have only shown results from two electrodes in a single experiment, as we did not have time for a thorough analysis of other similarly positioned electrodes from other, similar experiments. Such data is available from our collaborators at the Newcastle Uni-

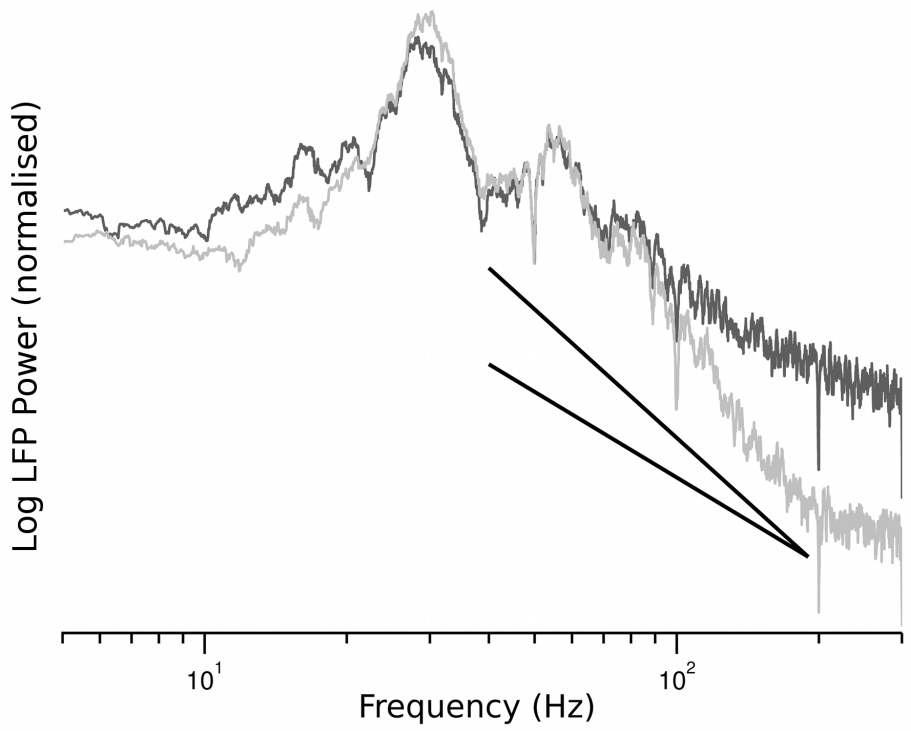

Figure 5.17: Example power spectra for LFPs recorded in macaque auditory neocortex with a Utah MEA (see Chapter 4), showing similarity of frequency-power scaling to the model results. The light grey line shows LFP power recorded in layer 1 (electrode 41), dark grey the LFP power recorded 400 microns below, in layer 2/3 (electrode 42). Scale bars show $1/f^2$ and $1/f^3$scaling. Increased power at higher frequencies compared to the model - particularly in layer 2/3 - is likely to be because of spikes, which are not present in the model LFP.

versity Institute of Neuroscience, so should can be explored further in conjunction with our model.

5.4 Summary

To investigate the LFP in more detail during gamma oscillations, we focussed on a simplified model of neocortical layer 2/3. We used this model to explore several aspects of LFP generation. Firstly, we found that perisomatic synaptic currents on pyramidal neurons resulting from basket interneuron firing dominate the LFP during gamma rhythms, in agreement with recent experimental results (Oren et al., 2010). We also predict that basket interneurons will contribute the majority of the LFP signal during random, uncorrelated activity because of the location of their synapses at pyramidal neuron somas. Secondly, we showed that an implementation of the proposed mechanism behind the experimentally observed phase-inversion in the LFP between layers 3 and 1 - a separation of the current source at the soma level and current sink a the apical dendrite level - matched experimental observations. Thirdly, we investigated how gamma power from a localised gamma-generating region in the model spread due to volume conduction. We found that the spatial range of the oscillation increased above and below the top of the soma-layer, and depended on the level of surrounding activity in the non-oscillating network. We also showed how the relative gamma frequency power above and below the point of phase inversion changes with soma depth. Finally, we investigated frequency scaling in the LFP power spectrum across space. We showed that LFPs at the apical dendrite level are predicted to echibit a low-pass filtering effect, which is absent from LFPs recorded at the soma level, due the relative dominance and localisation of the basket interneuron input. We confirmed that this matched the frequency scaling in our comparison experimental data, and, based on the frequency characteristics in our model, proposed that the experimentally recorded LFP could be contaminated by spiking activity at frequencies as low as $\sim$90 Hz.

While we have made some significant progress in understanding LFP generation using network models, there is always more work to be done. We have not yet created a model that explains the lateral patchy gamma power observed in experimental MEA recordings *in vitro*. An obvious extension to our modelling approach described

here would be to incorporate patchy connectivity into the pyramidal neuron's connection patterns using the models described by Voges et al. (Voges et al., 2009, 2010). The question of the frequency scaling properties of the LFP needs to be more fully addressed, which will involve further modelling work informed by experimental data. Investigation of network states other than gamma oscillations and random activity, and an extension to the *in vivo* setting, are also necessary.

In the next chapter we summarise the past four chapters and discuss our results. Finally, we propose various areas of research that could benefit from simulations of LFPs in spatially organised network models.

Chapter 6

Discussion & Outlook

The local field potential is a widely used measurement of neuronal network dynamics, yet it is still unclear exactly how the LFP relates to the underlying neuronal network activity (Einevoll et al., 2013). A better understanding of this relationship will improve our knowledge of neuronal network function and dysfunction by improving our ability to interpret experimental results. Recent modelling studies have begun to investigate the origin of the LFP and its relationship to neuronal activity (Bédard et al., 2006a; Bédard and Destexhe, 2009; Bédard et al., 2004; Diwakar et al., 2011; Lindén et al., 2010, 2011; Reimann et al., 2013), but so far these models have tended to examine LFPs created by specified synaptic inputs rather than generated by intrinsic network dynamics. We have attempted to add to these previous efforts by developing a simulation tool, using it to implement network models, and investigating the LFP generated across space in these models.

In this final chapter, we first summarise the key results of our research. We then describe how these results can be helpful for interpreting experimental recordings, as well as constraining network models. Finally we make suggestions for improvements to our methods and future research directions.

6.1 Summary of key results

6.1.1 Simplified compartmental model for LFP generation

Spiking neural network models are increasingly being used to investigate neuronal network dynamics, but there is currently no established method for reliably simu-

lating local field potentials resulting from activity in this kind of model (Einevoll et al., 2013). Various approximations to the LFP have been used when comparing model output with experimental recordings (beim Graben and Rodrigues, 2013; Mazzoni et al., 2008, 2010), but these approximations are not valid generally, and are particularly unsuitable for simulating the LFP in space - an important missing dimension if we are to understand the function of real neural networks. The physics underlying LFP generation are well established, though, and previous studies incorporating this knowledge have shown that the LFP is dependent on the distribution of current sources and sinks along a neuron's dendrites (Łęski et al., 2013; Lindén et al., 2010, 2011). This distribution can be calculated using compartmental neuron models. To ensure a realistic distribution of currents in space, these models used reconstructed neuron models containing hundreds of compartments. These are computationally expensive to simulate, especially in a network. We therefore tested reduced compartmental models containing $<$10 compartments (Bush and Sejnowski, 1993) to assess their suitability for modelling LFPs, with a view to using these less resource-intensive models in network simulations.

The reduced models were created from specific morphological reconstructions, so we compared simulated LFPs from populations of the reduced models with populations of the equivalent morphologically reconstructed models. We also compared LFPs simulated from populations of the same neuron type but from different whole cell reconstructions, to gauge the magnitude of the errors introduced by the reduction method in terms of the general variability between different neuron models. We found that the reduced neuron models reproduced the LFP magnitude, spatial range, and frequency scaling characteristics of the full cell reconstructions to a good degree of accuracy; the errors were of a similar size to the differences in these values observed when using the alternative full cell reconstructions. Care should be taken when interpreting results from the reduced compartmental models: our comparisons for different synaptic inputs were not exhaustive, and the reductions may not be appropriate when modelling active conductances. However, used appropriately, the reduced models will hasten the progress of LFP modelling research. We proceeded to use the reduced models for network simulations, but first we needed to build a simulator.

6.1.2 VERTEX simulation tool

Currently, one simulation tool is available for simulating extracellular potentials due to neuronal membrane currents (Lindén et al., 2014). LFPy allows the simulation of the extracellular potential generated by a single neuron, but does not currently allow network simulations. It is possible to generate network activity in a separate network simulator, and then simulate the LFP in a population one neuron at a time using this pre-generated spiking input (Lindén et al., 2011). However, simulating large numbers of neurons individually like this takes a long time. It also requires separate simulators to generate the spike times and simulate the LFPs. We developed the first tool designed to simulate the spatial characteristics of LFPs resulting from neuronal network activity: VERTEX. VERTEX is written in Matlab, and provides a simple user interface for specifying models, maximising its utility to the general neuroscience community. It also provides a significant speed improvement over LFPy when simulating large numbers of neurons. Many labs already use their own Matlab code for analysis of electrophysiology experiments. These tools can be applied directly to the data generated by VERTEX simulations, allowing users to run models and analyse their output from within the same computing environment. VERTEX's functionality is currently aimed particularly at exploring experiments *in vitro*, providing functions for calculating missing connectivity due to slice preparation, but networks exploring *in vivo* LFP generation can also be developed. Finally, VERTEX can be run in parallel using the Matlab Parallel Computing Toolbox, improving computation time approximately linearly in proportion to the number of parallel processes used. Though we have yet to test it on a cluster or in the cloud, running VERTEX in these distributed environments using the Matlab Distributed Computing Server is theoretically possible without modifying any of the code. This will allow the rapid simulation of large-scale models, or speed-up parameter exploration in smaller models. VERTEX's ability to scale to large-scale simulations was demonstrated by the implementation of a neocortical slice model containing 175 000 neurons.

6.1.3 Neocortical slice model

One of our goals for the development of VERTEX was to create a model that could be used in conjunction with recent recordings made by our collaborators using Utah multi-electrode arrays *in vitro*. The Utah array records across the depth and breadth of the slice, allowing the analysis of network dynamics across space both within and between cortical layers. We therefore created a neocortical slice model based on experimental anatomical and connectivity data, arranged in space to emulate the experimental setup. We stimulated the model to create a gamma oscillation, as observed experimentally after bath kainate application (Ainsworth et al., 2011; Cunningham et al., 2003; Fisahn et al., 1998). The neurons in the model were stimulated so that their spiking matched previously collected experimental data on neuron firing during persistent gamma oscillations. Given this spiking activity and the model connectivity, we used VERTEX to simulate the resultant LFPs measured by a simulated Utah array. The model reproduced the phase-inversion in the LFP observed between layer 3 and layer 1 in the experimental data, but the lateral coherence of the simulated LFP was much greater than in the experimental recordings. Variations in the gamma frequency power could result from many factors: local variations in connectivity, synaptic strengths, neuronal density, or individual neuron properties are all potential explanations. Given the currently available data, we suggest that the patchy connectivity of pyramidal neurons in layer 2/3 (Binzegger et al., 2007; Kisvárday and Eysel, 1992) could potentially result in these kinds of local variations in gamma power.

6.1.4 LFP features during network activity

We turned to a simpler model to investigate the generation of the LFP during gamma oscillations in more detail. First, we examined a highly simplified model of how a presynaptic neuron contributes to the LFP based on its synaptic connectivity and how it fires in relation to other neurons of the same type. This model allows us to predict theoretical upper and lower bounds for the relative magnitude of the LFP caused by different presynaptic neuron types, ignoring the effect of synaptic localisation on the dendrites. We next looked at simulations of localised gamma oscillations surrounded by randomly firing populations, to study the differences between these

two activity regimes and how the gamma oscillation signal spread due to volume conduction. We found that basket interneurons totally dominated the LFP in the gamma-oscillating region, but also made the largest contribution to the LFP during random spiking activity due to the localisation of their synapses on pyramidal neuron somas. The visibility of the gamma oscillation was smallest at the phase-inversion point and widened above and below this point, similar to the LFP range calculations in chapter 2, and those reported in (Lindén et al., 2011). The gamma coherence dropped abruptly outside of the gamma generating region; lower random activity did not change the abruptness of the coherence drop, but did extend the distance at which it occurred.

These results are important for the interpretation of experimental recordings. Firstly, we can see that gamma power does not spread far outside of the gamma generating region through volume conduction, given standard rates of neuronal firing outside of this region. We can be fairly sure, then, that high coherence between electrodes in an MEA indicates that there is network synchrony between those two points and not just that the oscillation is visible at both points due to volume conduction. Secondly, the dominance of basket interneuron firing on the LFP even during entirely random activity has interesting implications for the interpretation of LFPs in general. If LFPs primarily contain contributions from basket cells, then disorders known to affect basket cell activity or numbers should have measurable effects on the LFP across many states and not just during gamma oscillations. Further work is required to establish the general validity of this result: the simplified layer 2/3 model does not contain any synapses from external sources (other layers or thalamic inputs), which under different network conditions may contribute more to the LFP than the basket cells. However, it is to be in agreement with recent experimental observations regarding the visibility of inhibitory contributions to the LFP (Menendez de la Prida and Trevelyan, 2011; Oren et al., 2010; Trevelyan, 2009).

6.2 How can LFP simulations lead to a better understanding of neuronal dynamics?

The development of MEAs, and the increasing availability of cheap hardware to store the large data sets collected by MEAs, means that the LFP is likely to continue to

be used as an experimental measure of network activity for the foreseeable future. To make the best use of these measurements, it is essential to understand how they relate to the underlying neuronal activity. We will now consider how our results impact on the understanding of neuronal network activity, how the models and tools we developed can be used in future experiments to gain specific knowledge that would not be possible to obtain without simulating LFPs generated by neuronal network models.

6.2.1 Consideration of electrode position

We saw in Chapter 5 that the depth of the extracellular electrode, and thickness of the cortical layer, influences both the power and phase of the recorded LFP. An electrode positioned at the phase-inversion point will record a very low-power signal compared with an electrode a mere tens of micrometers above or below it. This has important implications for experiments comparing LFP characteristics recorded from two or more different groups of animals, where the groups have been subjected to different manipulation (gene knockout, for example). In these cases, it is important to check that cortical thickness, or the relative thickness of each cortical layer, has not also been affected. It is possible that some manipulations will lead to subtle differences in cortical thickness that result in electrodes being positioned at a different relative cortical depth between groups. Differences in LFP measures could then be a result of this positional difference, rather than a true reflection of altered network dynamics. A related effect has been demonstrated at the level of the EEG: the position of the brain - and therefore the distribution of cerebro-spinal fluid in the skull - depends on head angle, which can significantly affect signal power in different frequency bands (Rice et al., 2013).

6.2.2 Checking spiking models

Neuronal spiking is profoundly under-sampled in any electrophysiology experiment, even with the use of high density MEAs. Experimentally recorded LFPs are useful for gauging overall network activity at a higher level than looking at individual neuron spiking. Modelling the LFP resulting from what is thought to be representative spiking activity underlying a network state, as we have done for gamma oscilla-

tions (Chapter 4), allows us to check the consequences of this under-sampling: given the presumed pattern of spiking and the network connectivity, is the experimentally measured LFP reproduced by the model? If not, has there been a systematic bias in the spike sampling, or is the connectivity or spatial positioning model inadequate? We simulated a network using experimentally derived connectivity statistics, and tuned the model to reproduce the experimentally observed spiking patterns. However, this model did not account for the lateral variations in gamma frequency power that were observed experimentally. While this variation could have various different underlying causes, we suspect that deficiencies in our spatial connectivity model are the most likely explanation, as we did not include details of the patchy connectivity seen in the superficial layers of macaque and cat neocortex (Binzegger et al., 2007; Kisvárday and Eysel, 1992). Dynamics in networks with "patchy" connectivity have only just begun to be investigated (Bauer et al., 2012; Voges et al., 2009, 2010), but an obvious next step for our neocortical slice model would be to incorporate this spatial connectivity feature. Given the other potential explanations for this gamma power distribution, further experiments will also be required to verify whether patchy connectivity is the main contributor to the patchiness in gamma power. These experiments would, ideally, map out the gamma frequency power over the space of the slice, and subsequently co-register these maps with the patterns of synaptic boutons from layer 2/3 pyramidal neurons.

6.2.3 Verifying spike-field coherence measures

The LFP and MUA are usually treated as resulting from different processes when analysing neuronal activity. However, the separation between MUA and LFP is made by filtering at an arbitrary cutoff frequency, below which it is assumed synaptic activity is the dominant contributor and above which spiking is the dominant contributor. Several recent studies have shown that the LFP can contain a significant spiking component even at frequencies lower than 100 Hz (Schomburg et al., 2012; Zanos et al., 2011). A recent study systematically investigated the possible contamination of spikes into lower frequencies of the LFP, showing that, neurons firing at high rates with large spike-widths, the frequency contamination could theoretically extend as far down as 10 Hz (Waldert et al., 2013).

This contamination is important for the interpretation of simultaneously meas-

ured spiking and LFPs. The spike field coherence (SFC) is a measure of the phase synchronisation between the LFP and spikes, as a function of frequency (Fries et al., 2001). It indicates how strongly entrained a particular neuron is to a particular LFP frequency. However, if spikes contribute significantly to the LFP, then the SFC is really measuring how strongly entrained the neuron's spiking is to itself. Strategies for mitigating against this are discussed in (Waldert et al., 2013). We propose that network simulations of the LFP as detailed in this thesis can also be used to check experimental SFC results. LFPs can be simulated to only incorporate synaptic and return currents, as in Chapter 5. By implementing various network models that are thought to describe the experimental system under investigation and simulating spiking and the resultant "pure" LFPs in the network, a range of SFC predictions for a given experiment can be generated. The experimentally measured SFC can be compared to these predictions, and if it is within the simulated bounds, the experimenter can be reasonably sure that the SFC measurement is not contaminated by the low frequency component of the spikes.

6.2.4 Spatio-temporal dynamics of LFPs in disease models

Building a reliable and informative model of a disease process to explain an alteration in observed brain dynamics is a particularly difficult task, as a wide variety of different changes can lead to similar pathological network activity (Spencer, 2009a; Staley et al., 2011). We propose that simulating the LFP at different spatial locations in neuronal network models of disease could provide an extra constraint on such models, if relevant extracellular recordings are available for comparison. We consider two examples related to gamma rhythms and basket interneurons.

Basket interneurons are known to be crucial for the generation of gamma oscillations (Bartos et al., 2007), and in our investigations were shown to be the cells whose synaptic activity dominates the LFP during gamma rhythms. Basket interneurons are the primary parvalbumin (PV - a calcium buffering protein) expressing neuron in the cortex (Freund and Katona, 2007; Kawaguchi and Kubota, 1998; Kosaka et al., 1987; Zaitsev et al., 2005). It may come as little surprise, then, that disorders in which gamma rhythm alterations have been observed show changes in PV expression compared to healthy controls. Reduction in PV expression has been observed in tissue samples both from patients with several kinds of epilepsy Alonso-Nanclares

et al. (2005); Andrioli et al. (2007); DeFelipe et al. (1993); Marco and DeFelipe (1997), and patients with schizophrenia Beasley and Reynolds (1997); Behrens and Sejnowski (2009); Cunningham et al. (2006); Lewis et al. (2012). However, the phenotypes of these disorders are very different, and the changes in gamma frequency activity opposite (in general): in schizophrenia, gamma power is frequently observed to be reduced, or gamma synchrony between regions reduced (Gallinat et al., 2004; Green et al., 2003; Haig et al., 2000; Kissler et al., 2000; Light et al., 2006), whereas in some epilepsies gamma power increases have been observed interictally in generalised epilepsy (Willoughby et al., 2003) and prior to seizure onset in some focal epilepsies (Alarcon et al., 1995; Fisher et al., 1992).

It is still unclear as to whether the loss of PV in these conditions reflects the death of basket cells, or whether basket cells are still present but with reduced PV expression, or if there is a combination of these effects. Recent evidence from a rodent schizophrenia model suggested that basket interneurons are still present, but that a proportion of them do not express PV (Powell et al., 2012). A computational model implemented to investigate the dynamical consequences of various anatomical changes observed in schizophrenia found that reducing basket interneuron synaptic efficacy, or reducing the number of basket interneuron synapses - similar to reducing the number of basket interneurons through cell death - both reduced gamma frequency power (Spencer, 2009b). Another computational model of network gamma oscillations incorporating proposed effects of reduced PV expression at the basket interneuron synapses predicted a reduction in gamma power through increased asynchronous GABA release by the interneurons (Volman et al., 2011). If suitable MEA recordings from animal models of schizophrenia are available, they could be compared with simulated LFPs from spatially organised networks implementing each of these potential mechanisms, as each could produce a different *spatial* gamma oscillation profile. A most likely underlying mechanism for the observed changes could then be proposed and inform future experimental work.

Reductions in PV are also seen in epilepsy (Andrioli et al., 2007; DeFelipe et al., 1993; Marco and DeFelipe, 1997), that could also result from these cells ceasing PV expression (Scotti et al., 1997) or from cell death or reduction in connectivity (Marco et al., 1997; Marco and DeFelipe, 1997). However, the reductions in PV appear to be highly localised to abnormal regions. Experiments on resected human epileptic

neocortex by our collaborators at the Newcastle University Institute of Neuroscience have begun to investigate correlations between these localised anatomical changes and MEA recordings (Simon et al., 2014). Modelling the network in space will aid in our understanding of how abnormal activity can spread out of the abnormal anatomical regions, and simulating the LFP will allow model predictions to be compared directly with the experiments. We have already begun to add functionality to VERTEX to simplify the specification of spatially inhomogeneous networks; see Figure 6.1.

6.3 Avenues for future research

In addition to the nonspecific future research directions outlined above, we have several suggestions for specific projects that draw on our methods and results.

6.3.1 Improvements to VERTEX

Several improvements can be made to VERTEX regarding performance and functionality. Performance-wise, some optimisations are possible that have not yet been implemented through lack of time. Memory consumption was reduced to an extent by using the smallest possible data types for each kind of value rather than Matlab's standard 64-bit floating point representations. However, further memory savings can be made by incorporating the indexing scheme outlined in (Morrison et al., 2005). This allows small integers to be used to store indices: relative rather than absolute indices are used. The slowest running sections of code can also be re-implemented in C or Fortran to be called from Matlab using its MEX interface. In particular, the spike queue and delivery mechanism is only vectorised over spikes, so increasing spike rates increase code interpretation overheads. The spike queuing code is therefore a priority target for implementation in a compiled language. Finally, VERTEX should be thoroughly tested for use on clusters or cloud-based systems such as Amazon's Elastic Compute Cloud. Theoretically, VERTEX should work on any cluster running the Matlab Distributed Computing Server.

Regarding features, an important next step will be to allow users to import their own compartmental models rather than relying on the currently implemented neuron morphologies, or having to specify their own morphologies by hand. The ideal solu-

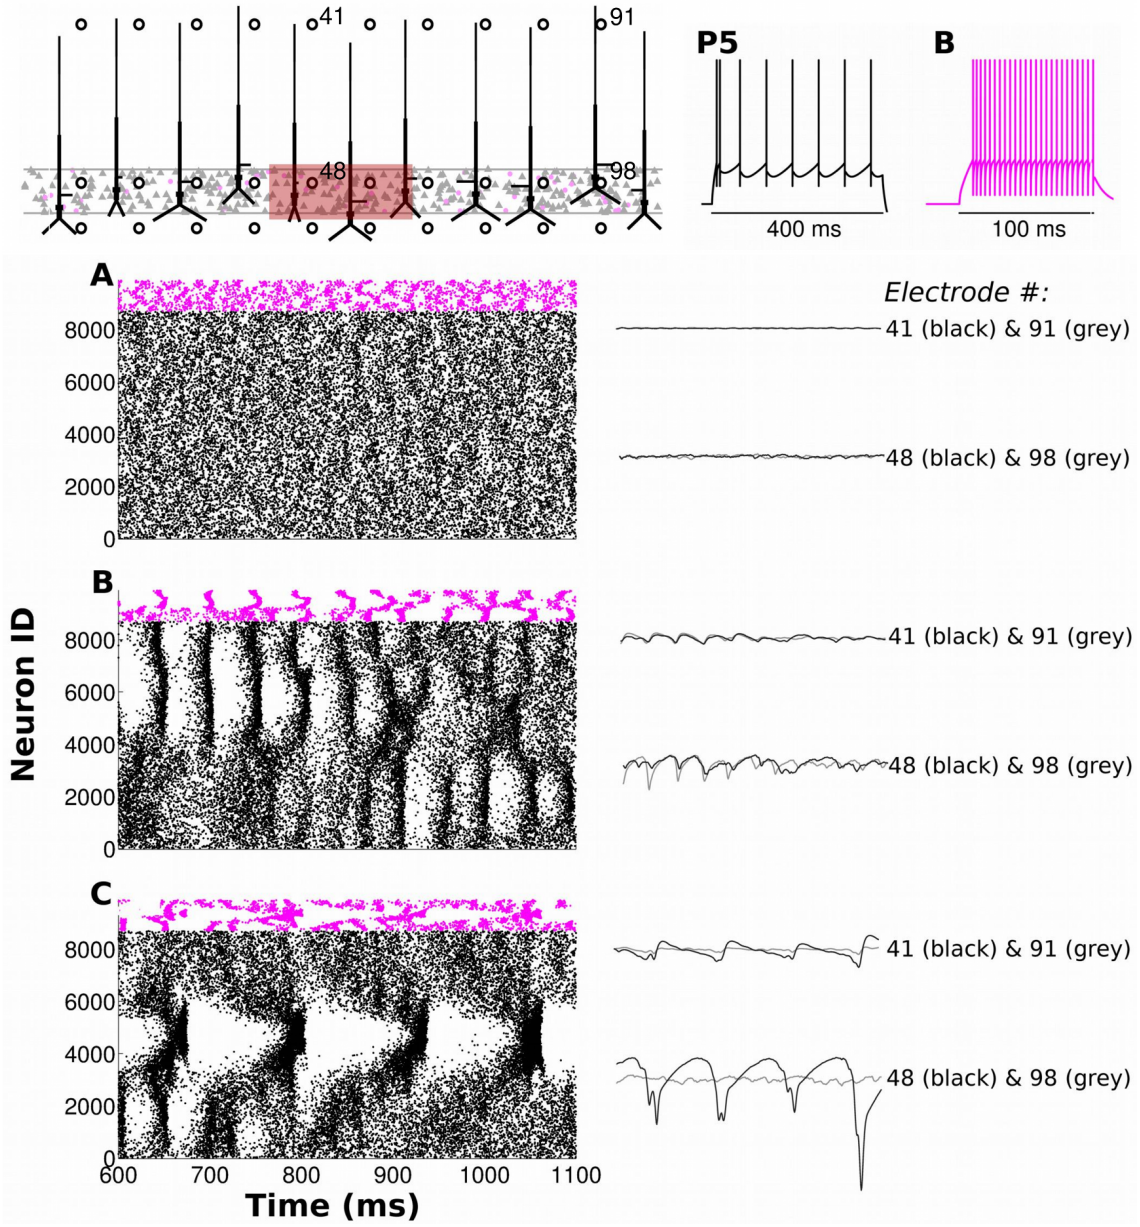

Figure 6.1: Example results from a model with inhomogeneous connectivity. Such models could be used to investigate spatially localised changes in anatomy as seen, for example, in focal epilepsy. Top: model structure and neuron dynamics. The model is 4 mm wide, and contains layer 5 pyramidal (P5) and layer 5 basket cells (B5) only. The red shaded region (1 mm wide) shows the localised region we modify the connectivity in to produce the bottom spike raster. (A) show the basline spiking activity and LFPs at different electrodes. (B) shows the same model with the number of excitatory synapses increased (simulating synaptic sprouting). (C) shows the same model, with the number of inhibitory synapses in the central region reduced by 80%.

tion would be to allow VERTEX to import standardised descriptions of full network models from (Gleeson et al., 2010) or to add compatibility with the PyNN Python interface (Davison et al., 2009); in the short term, an importer for single neuron morphologies from NEURON hoc files or NeuroML would be more practical, and still very useful. VERTEX can currently only specify random connectivity based on spatial distributions. It might be useful to import specific connectivity patterns that have been investigated separately, or have been collected from experimental data as appears to be increasingly likely over the next few years. Adding further, more elaborate - but still probabilistic - connectivity patterns will also aid with the investigation of network synchrony over space. The superficial neocortex is known to contain patchy projections from pyramidal neurons, a structure that is not well described by the currently implemented uniform or Gaussian spatial decay probabilities. This may help to answer questions on the role of connectivity in synchronous activity across the network. Finally, it will be important in future to allow users to specify arbitrary model shapes. Cortical tissue is highly convoluted, and large MEAs can extend over regions that include curved tissue. As pyramidal neuron membrane currents dominate the LFP, and they are aligned with their apical dendrites perpendicular to the cortical surface, the LFP across space will be dependent on this curvature. This will be important to take into account for future comparisons with experimental MEA recordings.

6.3.2 Closed-loop stimulation modelling

Recent advances in intra-cranial implants have led to the development of closed-loop stimulation devices for epilepsy (Berényi et al., 2012; Fountas and Smith, 2007; Smith et al., 2010). The term "closed-loop" refers to the manner in which the signal measured from the brain drives the stimulating device, and this stimulation then alters the brain activity, and thus the measured signal, and so on. This is in contrast to more conventional brain stimulation systems such as the Deep Brain Stimulation (DBS) devices used, for example, to control the symptoms of Parkinson's disease (Benabid et al., 2000), in which stimulation is constantly applied, or switched on and off by the physician or patient rather than stimulating in response to measured activity. Closed loop systems are particularly important for managing seizures, as the stimulation must be applied rapidly in response to the first sign of a seizure.

The problem to be solved for closed loop stimulation is how best to stimulate to reduce the likelihood of a seizure, given the measured activity.

Modelling can make a significant contribution to solving this problem. In currently available closed-loop systems such as that described in (Smith et al., 2010), measurement and stimulation are performed by electrodes. Modelling the network activity believed to cause the LFP that triggers stimulation in this system will allow different stimulation scenarios to be tested, so that the best simulated outcomes can be trialled in the real system. This kind of simulation would need to take into account the location of the measurement electrodes in relation to the neurons and how the errant behaviour that the stimulator looks for comes about. This kind of simulation is possible in VERTEX. Ideally, we would also like to close the loop in the simulation, so that we could predict, given a particular pattern of potentially pathological activity, the optimal timing, location and magnitude of the stimulation. Extracellular stimulating electrodes are not yet implemented in VERTEX; this is certainly a feature that should be added. However, we have already used VERTEX in preliminary tests of a proposed closed-loop stimulation system that targets a specific subset of optogenetically altered neurons using light (Andrew Jackson, personal communication). In the proposed system, a subset of neurons is altered using a safe viral vector to express channelrhodopsins in their cell membranes. The extracellular potential is continuously monitored and modulates the strength of a light source. The benefit of this system is that the stimulation can be targeted to a specific subset of neurons, rather than being applied generally to the network via extracellular current injection.

Figure 6.2 shows an example preliminary attempt at modelling this kind of closed-loop feedback setup. We used a simplified model of neocortical layer 5, containing 8500 pyramidal neurons and 1500 basket interneurons arranged in a cylinder of radius 800 microns and soma-depth of 290 microns, with similar connectivity statistics to layer 5 in our neocortical slice model detailed in Chapter 4. We measured the LFP at the apical dendrite level, at the centre of the cylinder. We applied random currents to the neurons to simulate a persistent gamma oscillation, as described in the preceding chapters. After 1000 ms, we switched on the feedback stimulation. This consisted of applying a current proportional to the half-wave rectified LFP signal to 10% of the pyramidal neurons in the population, representing the opto-

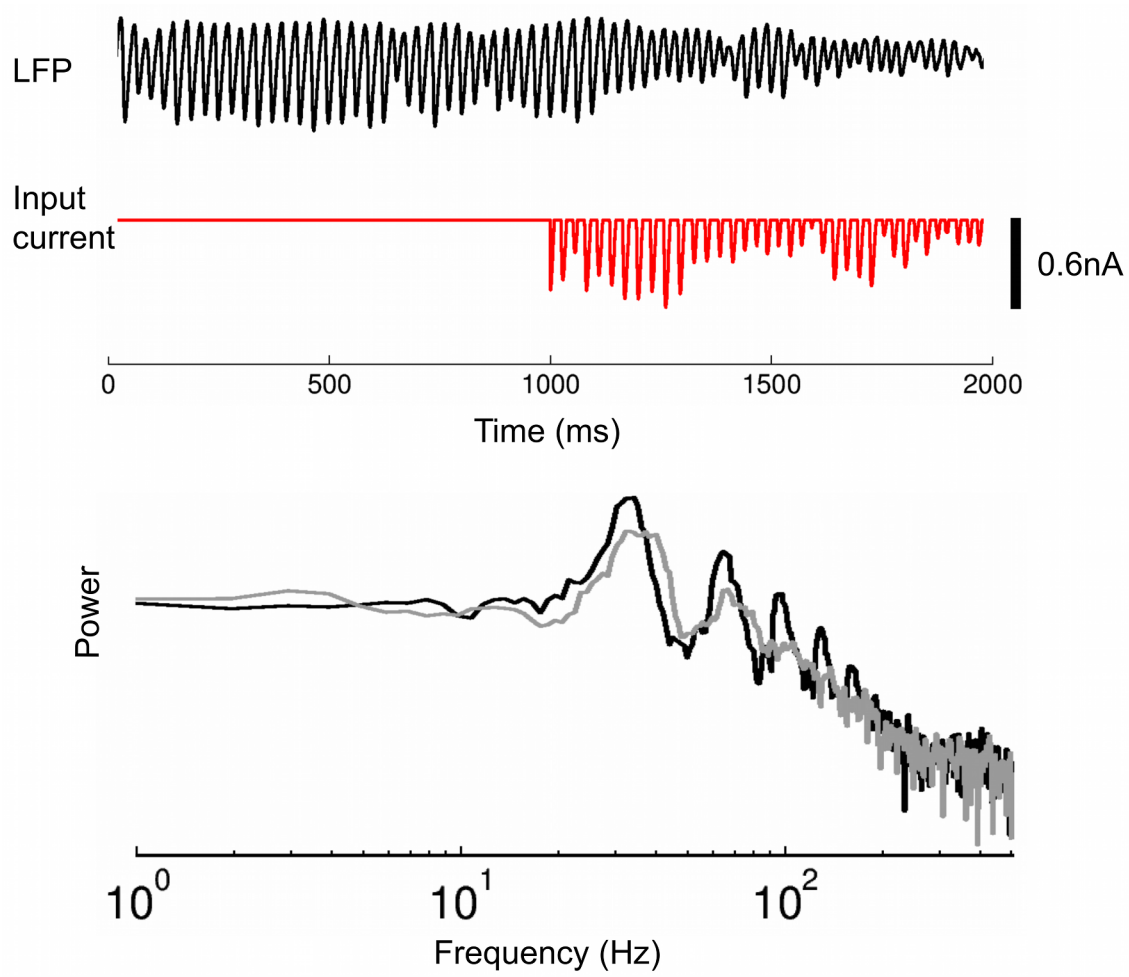

Figure 6.2: Model of closed-loop stimulation to reduce gamma oscillation power. Top: LFP recorded from a simulation of a gamma oscillation in neocortical layer 5. After 1 s, a "stimulator" is switched on. This simulates the currents in a subset (10%) of the pyramidal cell population created by a light source whose intensity is controlled by the half-wave rectified LFP signal. Bottom: Reduction in gamma power due to stimulation. The black line shows the LFP power spectrum prior to stimulation, and the grey line shows the LFP power spectrum during stimulation.

genetically altered neurons. As can be seen in Figure 6.2, the power of the gamma oscillation decreases, confirmed in the LFP power spectra calculated both before and during feedback stimulation (Figure 6.2, bottom).

While this model makes many simplifications that require further investigation (the kinetics of channelrhodopsin, for instance), it shows how simulations of the LFP from network models already have an important role to play in the development of new treatments. This kind of model will help to direct investigations into:

- which subgroup of neurons is best to stimulate, and at what time, given the LFP,
- how large this stimulation should be to reduce the likelihood of a seizure but have minimal other effects on the network,
- whether stimulating different sub-populations with different channelrhodopsins could result in improved control of the network (Carter and de Lecea, 2011; LaLumiere, 2011),
- how electrode position affects the LFP (this is very important if, for example, we want to stimulate during peaks or troughs of an oscillation whose phase is dependent on electrode position, as in gamma rhythms).

VERTEX is ideally suited to implementing models to investigate these questions.

6.3.3 Ephaptic coupling

Several exciting experimental results have recently shown that neuronally generated electric fields impact on the membrane potentials of nearby neurons without requiring any synaptic contact. Such "ephaptic" coupling of neurons was investigated in a modelling study in the late 1990s (Holt and Koch, 1999), and more recently confirmed in experiments showing that such interactions could entrain action potentials (Anastassiou et al., 2011) and potentially contribute to the spread of epileptiform activity (Zhang et al., 2014). We have purposefully ignored the contribution of ephaptic interactions in our models for the sake of simplicity. However, the recent experimental results show that the role of ephaptic interactions on network activity must be investigated further. As VERTEX can simulate the LFP at arbitrary locations in a network, it would be possible to incorporate an ephaptic

coupling mechanism that depended on the LFP. However, doing this rigorously would entail measuring the LFP near every compartment in the model, which is not feasible. Developing suitable approximation methods for incorporating realistic ephaptic coupling is therefore an important direction for future research.

Bibliography

Abeles, M. (1991). *Corticonics: Neural Circuits of the Cerebral Cortex*. Cambridge University Press.

Adrian, E. D. and Moruzzi, G. (1939). Impulses in the pyramidal tract. *The Journal of Physiology*, 97(2):153–199. PMID: 16995153 PMCID: PMC1393899.

Ahlfors, S. P., Han, J., Belliveau, J. W., and Hämäläinen, M. S. (2010). Sensitivity of MEG and EEG to source orientation. *Brain Topography*, 23(3):227–232.

Ainsworth, M. (2013). *Cross species comparison of the spatiotemporal properties of the gamma frequency oscillation*. PhD thesis, Newcastle University, Newcastle upon Tyne.

Ainsworth, M., Lee, S., Cunningham, M. O., Roopun, A. K., Traub, R. D., Kopell, N. J., and Whittington, M. A. (2011). Dual gamma rhythm generators control interlaminar synchrony in auditory cortex. *The Journal of Neuroscience*, 31(47):17040–17051. PMID: 22114273 PMCID: PMC3654396.

Ainsworth, M., Lee, S., Cunningham, M. O., Traub, R. D., Kopell, N. J., and Whittington, M. A. (2012). Rates and rhythms: A synergistic view of frequency and temporal coding in neuronal networks. *Neuron*, 75(4):572–583.

Alarcon, G., Binnie, C., Elwes, R., and Polkey, C. (1995). Power spectrum and intracranial EEG patterns at seizure onset in partial epilepsy. *Electroencephalography and Clinical Neurophysiology*, 94(5):326–337.

Alonso-Nanclares, L., Garbelli, R., Sola, R., Pastor, J., Tassi, L., Spreafico, R., and DeFelipe, J. (2005). Microanatomy of the dysplastic neocortex from epileptic patients. *Brain*, 128(1):158–173.

Anastassiou, C. A., Perin, R., Markram, H., and Koch, C. (2011). Ephaptic coupling of cortical neurons. *Nature Neuroscience*, 14(2):217–223.

Andrioli, A., Alonso-Nanclares, L., Arellano, J., and DeFelipe, J. (2007). Quantitative analysis of parvalbumin-immunoreactive cells in the human epileptic hippocampus. *Neuroscience*, 149(1):131–143.

Aoki, F., Fetz, E. E., Shupe, L., Lettich, E., and Ojemann, G. A. (1999). Increased gamma-range activity in human sensorimotor cortex during performance of visuomotor tasks. *Clinical Neurophysiology*, 110(3):524–537. PMID: 10363776.

Aronson, D. G., Ermentrout, G. B., and Kopell, N. (1990). Amplitude response of coupled oscillators. *Physica D: Nonlinear Phenomena*, 41(3):403–449.

Arsiero, M., Luscher, H.-R., Lundstrom, B. N., and Giugliano, M. (2007). The impact of input fluctuations on the frequency-current relationships of layer 5 pyramidal neurons in the rat medial prefrontal cortex. *The Journal of Neuroscience*, 27(12):3274–3284.

Ascoli, G. A., Donohue, D. E., and Halavi, M. (2007). NeuroMorpho.Org: a central resource for neuronal morphologies. *The Journal of Neuroscience*, 27(35):9247–9251. PMID: 17728438.

Atallah, B. V. and Scanziani, M. (2009). Instantaneous modulation of gamma oscillation frequency by balancing excitation with inhibition. *Neuron*, 62(4):566–577.

Badel, L., Lefort, S., Brette, R., Petersen, C. C. H., Gerstner, W., and Richardson, M. J. E. (2008). Dynamic i-v curves are reliable predictors of naturalistic pyramidal-neuron voltage traces. *Journal of Neurophysiology*, 99(2):656–666.

Ball, T., Nawrot, M. P., Pistohl, T., Aertsen, A., Schulze-bonhage, A., and Mehring, C. (2004). Towards an implantable brain-machine interface based on epicortical field potentials. In *Biomedizinische Technik, Biomedical Engineering, 38. Jahrestagung der Deutschen Gesellschaft für Biomedizinische Technik im VDE - BMT 2004*, page 756–759.

Bartos, M., Vida, I., and Jonas, P. (2007). Synaptic mechanisms of synchronized gamma oscillations in inhibitory interneuron networks. *Nature Reviews Neuroscience*, 8(1):45–56.

Başar, E. (2013). A review of gamma oscillations in healthy subjects and in cognitive impairment. *International Journal of Psychophysiology*.

Bauer, M., Oostenveld, R., Peeters, M., and Fries, P. (2006). Tactile spatial attention enhances gamma-band activity in somatosensory cortex and reduces low-frequency activity in parieto-occipital areas. *The Journal of Neuroscience*, 26(2):490–501. PMID: 16407546.

Bauer, R., Zubler, F., Hauri, A., Muir, D. R., and Douglas, R. J. (2012). Developmental origin of patchy axonal connectivity in the neocortex: A computational model. *Cerebral Cortex*, page bhs327. PMID: 23131803.

Bazelot, M., Dinocourt, C., Cohen, I., and Miles, R. (2010). Unitary inhibitory field potentials in the CA3 region of rat hippocampus. *The Journal of Physiology*, 588(Pt 12):2077–2090. PMID: 20403979 PMCID: PMC2911213.

Beasley, C. L. and Reynolds, G. P. (1997). Parvalbumin-immunoreactive neurons are reduced in the prefrontal cortex of schizophrenics. *Schizophrenia Research*, 24(3):349–355. PMID: 9134596.

Bédard, C. and Destexhe, A. (2009). Macroscopic models of local field potentials and the apparent 1/f noise in brain activity. *Biophysical Journal*, 96(7):2589–2603.

Bedard, C. and Destexhe, A. (2013). Reply to gratiy et al. *Journal of Neurophysiology*, 109(6):1683–1683. PMID: 23503557.

Bédard, C., Kröger, H., and Destexhe, A. (2004). Modeling extracellular field potentials and the frequency-filtering properties of extracellular space. *Biophysical Journal*, 86(3):1829–1842.

Bédard, C., Kröger, H., and Destexhe, A. (2006a). Does the 1/f frequency scaling of brain signals reflect self-organized critical states? *Physical Review Letters*, 97(11).

Bédard, C., Kröger, H., and Destexhe, A. (2006b). Model of low-pass filtering of local field potentials in brain tissue. *Physical Review E*, 73(5).

Bédard, C., Rodrigues, S., Roy, N., Contreras, D., and Destexhe, A. (2010). Evidence for frequency-dependent extracellular impedance from the transfer function

between extracellular and intracellular potentials: intracellular-LFP transfer function. *Journal of Computational Neuroscience*, 29(3):389–403. PMID: 20559865.

Behrens, M. M. and Sejnowski, T. J. (2009). Does schizophrenia arise from oxidative dysregulation of parvalbumin-interneurons in the developing cortex? *Neuropharmacology*, 57(3):193–200.

beim Graben, P. and Rodrigues, S. (2013). A biophysical observation model for field potentials of networks of leaky integrate-and-fire neurons. *Frontiers in Computational Neuroscience*, 6.

Benabid, A. L., Krack, P. P., Benazzouz, A., Limousin, P., Koudsie, A., and Pollak, P. (2000). Deep brain stimulation of the subthalamic nucleus for parkinson's disease: methodologic aspects and clinical criteria. *Neurology*, 55(12 Suppl 6):S40–44. PMID: 11188974.

Benedek, G. B. and Villars, F. M. H. (2000). *Physics With Illustrative Examples From Medicine and Biology: Statistical Physics*. Springer.

Berens, P., Keliris, G. A., Ecker, A. S., Logothetis, N. K., and Tolias, A. S. (2008). Comparing the feature selectivity of the gamma-band of the local field potential and the underlying spiking activity in primate visual cortex. *Frontiers in Systems Neuroscience*, 2:2. PMID: 18958246.

Berényi, A., Belluscio, M., Mao, D., and Buzsáki, G. (2012). Closed-loop control of epilepsy by transcranial electrical stimulation. *Science*, 337(6095):735–737. PMID: 22879515.

Binzegger, T., Douglas, R. J., and Martin, K. A. C. (2004). A quantitative map of the circuit of cat primary visual cortex. *The Journal of Neuroscience*, 24(39):8441–8453. PMID: 15456817.

Binzegger, T., Douglas, R. J., and Martin, K. A. C. (2007). Stereotypical bouton clustering of individual neurons in cat primary visual cortex. *The Journal of Neuroscience*, 27(45):12242–12254.

Blasdel, G. G., Lund, J. S., and Fitzpatrick, D. (1985). Intrinsic connections of macaque striate cortex: axonal projections of cells outside lamina 4C. *The Journal of Neuroscience*, 5(12):3350–3369.

Börgers, C., Epstein, S., and Kopell, N. J. (2005). Background gamma rhythmicity and attention in cortical local circuits: A computational study. *Proceedings of the National Academy of Sciences of the United States of America*, 102(19):7002–7007. PMID: 15870189.

Börgers, C. and Kopell, N. (2003). Synchronization in networks of excitatory and inhibitory neurons with sparse, random connectivity. *Neural Computation*, 15(3):509–538.

Börgers, C. and Kopell, N. (2005). Effects of noisy drive on rhythms in networks of excitatory and inhibitory neurons. *Neural Computation*, 17(3):557–608.

Böttger, D., Herrmann, C., and Von Cramon, D. (2002). Amplitude differences of evoked alpha and gamma oscillations in two different age groups. *International Journal of Psychophysiology*, 45(3):245–251.

Bower, J. M. and Beeman, D. (1995). *The book of GENESIS: exploring realistic neural models with the GEneral NEural SImulation System.* TELOS, Santa Clara, Calif.

Bragin, A., Jando, G., Nadasdy, Z., Hetke, J., Wise, K., and Buzsaki, G. (1995). Gamma (40-100 hz) oscillation in the hippocampus of the behaving rat. *The Journal of Neuroscience*, 15(1):47–60. PMID: 7823151.

Brette, R. and Destexhe, A., editors (2012). *Handbook of Neural Activity Measurement.* Cambridge University Press, New York, NY, USA.

Brette, R. and Gerstner, W. (2005). Adaptive exponential integrate-and-fire model as an effective description of neuronal activity. *Journal of Neurophysiology*, 94(5):3637–3642.

Brette, R. and Goodman, D. F. M. (2011). Vectorized algorithms for spiking neural network simulation. *Neural Computation*, 23(6):1503–1535.

Brette, R., Rudolph, M., Carnevale, T., Hines, M., Beeman, D., Bower, J. M., Diesmann, M., Morrison, A., Goodman, P. H., Harris, Frederick C, J., Zirpe, M., Natschläger, T., Pecevski, D., Ermentrout, B., Djurfeldt, M., Lansner, A., Rochel, O., Vieville, T., Muller, E., Davison, A. P., El Boustani, S., and Destexhe, A.

(2007). Simulation of networks of spiking neurons: a review of tools and strategies. *Journal of Computational Neuroscience*, 23(3):349–398. PMID: 17629781.

Brunel, N. (2000). Dynamics of sparsely connected networks of excitatory and inhibitory spiking neurons. *Journal of Computational Neuroscience*, 8(3):183–208. PMID: 10809012.

Brunel, N. and Hakim, V. (1999). Fast global oscillations in networks of integrate-and-fire neurons with low firing rates. *Neural computation*, 11(7):1621–1671. PMID: 10490941.

Brunel, N. and Hakim, V. (2008). Sparsely synchronized neuronal oscillations. *Chaos (Woodbury, N.Y.)*, 18(1):015113. PMID: 18377094.

Brunel, N. and Hansel, D. (2006). How noise affects the synchronization properties of recurrent networks of inhibitory neurons. *Neural computation*, 18(5):1066–1110. PMID: 16595058.

Brunel, N. and Wang, X.-J. (2003). What determines the frequency of fast network oscillations with irregular neural discharges? i. synaptic dynamics and excitation-inhibition balance. *Journal of Neurophysiology*, 90(1):415–430. PMID: 12611969.

Buhl, D. L., Harris, K. D., Hormuzdi, S. G., Monyer, H., and Buzsáki, G. (2003). Selective impairment of hippocampal gamma oscillations in connexin-36 knock-out mouse in vivo. *The Journal of Neuroscience*, 23(3):1013–1018. PMID: 12574431.

Buhl, E. H., Tamás, G., and Fisahn, A. (1998). Cholinergic activation and tonic excitation induce persistent gamma oscillations in mouse somatosensory cortex in vitro. *The Journal of Physiology*, 513(1):117–126. PMID: 9782163.

Burns, S. P., Xing, D., and Shapley, R. M. (2011). Is gamma-band activity in the local field potential of v1 cortex a "Clock" or filtered noise? *The Journal of Neuroscience*, 31(26):9658–9664.

Bush, P. C. and Sejnowski, T. J. (1993). Reduced compartmental models of neocortical pyramidal cells. *Journal of Neuroscience Methods*, 46(2):159–166.

Buzsáki, G. (2004). Large-scale recording of neuronal ensembles. *Nature Neuroscience*, 7(5):446–451.

Buzsáki, G. (2006). *Rhythms of the brain*. Oxford University Press, Oxford; New York.

Buzsáki, G., Anastassiou, C. A., and Koch, C. (2012). The origin of extracellular fields and currents - EEG, ECoG, LFP and spikes. *Nature Reviews Neuroscience*, 13(6):407–420.

Buzsáki, G. and Draguhn, A. (2004). Neuronal oscillations in cortical networks. *Science*, 304(5679):1926–1929. PMID: 15218136.

Callaway, E. M. (1998). Local circuits in primary visual cortex of the macaque monkey. *Annual Review of Neuroscience*, 21(1):47–74. PMID: 9530491.

Carnevale, N. T. and Hines, M. L. (2006). *The NEURON Book*. Cambridge University Press.

Carter, M. E. and de Lecea, L. (2011). Optogenetic investigation of neural circuits in vivo. *Trends in Molecular Medicine*, 17(4):197–206. PMID: 21353638 PMCID: PMC3148823.

Catafau, A. M., Searle, G. E., Bullich, S., Gunn, R. N., Rabiner, E. A., Herance, R., Radua, J., Farre, M., and Laruelle, M. (2010). Imaging cortical dopamine d1 receptors using [11C]NNC112 and ketanserin blockade of the 5-HT2A receptors. *Journal of Cerebral Blood Flow & Metabolism*, 30(5):985–993.

Chang, Y.-T., Chen, P.-C., Tsai, I.-J., Sung, F.-C., Chin, Z.-N., Kuo, H.-T., Tsai, C.-H., and Chou, I.-C. (2011). Bidirectional relation between schizophrenia and epilepsy: a population-based retrospective cohort study. *Epilepsia*, 52(11):2036–2042. PMID: 21929680.

Chemla, S. and Chavane, F. (2010a). A biophysical cortical column model to study the multi-component origin of the VSDI signal. *NeuroImage*, 53(2):420–438.

Chemla, S. and Chavane, F. (2010b). Voltage-sensitive dye imaging: Technique review and models. *Journal of Physiology - Paris*, 104(1–2):40–50.

Chemla, S. and Chavane, F. (2012). Voltage-sensitive dye imaging. In Brette, R. and Destexhe, A., editors, *Handbook of Neural Activity Measurement*, pages 92–135. Cambridge University Press, New York, NY, USA.

Chrobak, J. J. and Buzsáki, G. (1998). Gamma oscillations in the entorhinal cortex of the freely behaving rat. *The Journal of Neuroscience*, 18(1):388–398. PMID: 9412515.

Clopath, C., Jolivet, R., Rauch, A., Lüscher, H.-R., and Gerstner, W. (2007). Predicting neuronal activity with simple models of the threshold type: Adaptive exponential integrate-and-fire model with two compartments. *Neurocomputing*, 70(10-12):1668–1673.

Colgin, L. L., Denninger, T., Fyhn, M., Hafting, T., Bonnevie, T., Jensen, O., Moser, M.-B., and Moser, E. I. (2009). Frequency of gamma oscillations routes flow of information in the hippocampus. *Nature*, 462(7271):353–357.

Contreras, D., Destexhe, A., and Steriade, M. (1997). Intracellular and computational characterization of the intracortical inhibitory control of synchronized thalamic inputs in vivo. *Journal of Neurophysiology*, 78(1):335–350. PMID: 9242284.

Csicsvari, J., Jamieson, B., Wise, K. D., and Buzsáki, G. (2003). Mechanisms of gamma oscillations in the hippocampus of the behaving rat. *Neuron*, 37(2):311–322. PMID: 12546825.

Cunningham, M. O., Davies, C. H., Buhl, E. H., Kopell, N., and Whittington, M. A. (2003). Gamma oscillations induced by kainate receptor activation in the entorhinal cortex in vitro. *The Journal of Neuroscience*, 23(30):9761–9769. PMID: 14586003.

Cunningham, M. O., Halliday, D. M., Davies, C. H., Traub, R. D., Buhl, t. l. E. H., and Whittington, M. A. (2004a). Coexistence of gamma and high-frequency oscillations in rat medial entorhinal cortex in vitro. *The Journal of Physiology*, 559(2):347–353. PMID: 15254156.

Cunningham, M. O., Hunt, J., Middleton, S., LeBeau, F. E. N., Gillies, M. J., Gillies, M. G., Davies, C. H., Maycox, P. R., Whittington, M. A., and Racca, C. (2006). Region-specific reduction in entorhinal gamma oscillations and parvalbumin-immunoreactive neurons in animal models of psychiatric illness. *The Journal of Neuroscience*, 26(10):2767–2776. PMID: 16525056.

Cunningham, M. O., Whittington, M. A., Bibbig, A., Roopun, A., LeBeau, F. E., Vogt, A., Monyer, H., Buhl, E. H., and Traub, R. D. (2004b). A role for fast rhythmic bursting neurons in cortical gamma oscillations in vitro. *Proceedings of the National Academy of Sciences of the United States of America*, 101(18):7152–7157.

Darvas, F., Pantazis, D., Kucukaltun-Yildirim, E., and Leahy, R. M. (2004). Mapping human brain function with MEG and EEG: methods and validation. *NeuroImage*, 23, Supplement 1:S289–S299.

Davison, A. P., Brüderle, D., Eppler, J., Kremkow, J., Muller, E., Pecevski, D., Perrinet, L., and Yger, P. (2009). PyNN: a common interface for neuronal network simulators. *Frontiers in Neuroinformatics*, 2:11.

DeFelipe, J. (1997). Types of neurons, synaptic connections and chemical characteristics of cells immunoreactive for calbindin-D28K, parvalbumin and calretinin in the neocortex. *Journal of Chemical Neuroanatomy*, 14(1):1–19.

DeFelipe, J., López-Cruz, P. L., Benavides-Piccione, R., Bielza, C., Larrañaga, P., Anderson, S., Burkhalter, A., Cauli, B., Fairén, A., Feldmeyer, D., Fishell, G., Fitzpatrick, D., Freund, T. F., González-Burgos, G., Hestrin, S., Hill, S., Hof, P. R., Huang, J., Jones, E. G., Kawaguchi, Y., Kisvárday, Z., Kubota, Y., Lewis, D. A., Marín, O., Markram, H., McBain, C. J., Meyer, H. S., Monyer, H., Nelson, S. B., Rockland, K., Rossier, J., Rubenstein, J. L. R., Rudy, B., Scanziani, M., Shepherd, G. M., Sherwood, C. C., Staiger, J. F., Tamás, G., Thomson, A., Wang, Y., Yuste, R., and Ascoli, G. A. (2013). New insights into the classification and nomenclature of cortical GABAergic interneurons. *Nature Reviews Neuroscience*, 14(3):202–216.

DeFelipe, J., Sola, R. G., Marco, P., Rio, M. d. R. d., Pulido, P., and Cajal, S. R. y. (1993). Selective changes in the microorganization of the human epileptogenic neocortex revealed by parvalbumin immunoreactivity. *Cerebral Cortex*, 3(1):39–48. PMID: 7679938.

Dehghani, N., Bédard, C., Cash, S. S., Halgren, E., and Destexhe, A. (2010). Comparative power spectral analysis of simultaneous elecroencephalographic and mag-

netoencephalographic recordings in humans suggests non-resistive extracellular media. *Journal of Computational Neuroscience*, 29(3):405–421. PMID: 20697790.

Destexhe, A. (2001). Simplified models of neocortical pyramidal cells preserving somatodendritic voltage attenuation. *Neurocomputing*, 38:167–173.

Destexhe, A. and Bedard, C. (2013). Local field potential. *Scholarpedia*, 8(8):10713.

Destexhe, A., Rudolph, M., Fellous, J.-M., and Sejnowski, T. J. (2001). Fluctuating synaptic conductances recreate in vivo-like activity in neocortical neurons. *Neuroscience*, 107(1):13–24.

Dickson, C. T., Biella, G., and Curtis, M. d. (2000). Evidence for spatial modules mediated by temporal synchronization of carbachol-induced gamma rhythm in medial entorhinal cortex. *The Journal of Neuroscience*, 20(20):7846–7854. PMID: 11027250.

Diwakar, S., Lombardo, P., Solinas, S., Naldi, G., and D'Angelo, E. (2011). Local field potential modeling predicts dense activation in cerebellar granule cells clusters under LTP and LTD control. *PLOS ONE*, 6(7):e21928.

Draguhn, A., Traub, R. D., Schmitz, D., and Jefferys, J. G. R. (1998). Electrical coupling underlies high-frequency oscillations in the hippocampus in vitro. *Nature*, 394(6689):189–192.

Eckhorn, R., Bauer, R., Jordan, W., Brosch, M., Kruse, W., Munk, M., and Reitboeck, H. J. (1988). Coherent oscillations: A mechanism of feature linking in the visual cortex? *Biological Cybernetics*, 60(2):121–130.

Eckhorn, R., Reitboeck, H. J., Arndt, M., and Dicke, P. (1990). Feature linking via synchronization among distributed assemblies: Simulations of results from cat visual cortex. *Neural Computation*, 2(3):293–307.

Economo, M. N. and White, J. A. (2012). Membrane properties and the balance between excitation and inhibition control gamma-frequency oscillations arising from feedback inhibition. *PLOS Computational Biology*, 8(1):e1002354.

Einevoll, G. T., Franke, F., Hagen, E., Pouzat, C., and Harris, K. D. (2012). Towards reliable spike-train recordings from thousands of neurons with multielectrodes. *Current Opinion in Neurobiology*, 22(1):11–17. PMID: 22023727 PMCID: PMC3314330.

Einevoll, G. T., Kayser, C., Logothetis, N. K., and Panzeri, S. (2013). Modelling and analysis of local field potentials for studying the function of cortical circuits. *Nature Reviews Neuroscience*, 14(11):770–785.

Einevoll, G. T., Pettersen, K. H., Devor, A., Ulbert, I., Halgren, E., and Dale, A. M. (2007). Laminar population analysis: Estimating firing rates and evoked synaptic activity from multielectrode recordings in rat barrel cortex. *Journal of Neurophysiology*, 97(3):2174–2190.

Engel, A. K., Fries, P., and Singer, W. (2001). Dynamic predictions: Oscillations and synchrony in top–down processing. *Nature Reviews Neuroscience*, 2(10):704–716.

Engel, A. K., Kreiter, A. K., König, P., and Singer, W. (1991). Synchronization of oscillatory neuronal responses between striate and extrastriate visual cortical areas of the cat. *Proceedings of the National Academy of Sciences*, 88(14):6048–6052. PMID: 2068083.

Ferezou, I., Bolea, S., and Petersen, C. C. H. (2006). Visualizing the cortical representation of whisker touch: Voltage-sensitive dye imaging in freely moving mice. *Neuron*, 50(4):617–629.

Ferguson, K. A., Williams, S., and Skinner, F. K. (2013). Experimentally constrained CA1 fast-firing parvalbumin-positive interneuron network models exhibit sharp transitions into coherent high frequency rhythms. *Frontiers in Computational Neuroscience*, 7:144.

Fisahn, A., Pike, F. G., Buhl, E. H., and Paulsen, O. (1998). Cholinergic induction of network oscillations at 40 hz in the hippocampus in vitro. *Nature*, 394(6689):186–189.

Fisher, R., Webber, W., Lesser, R., Arroyo, S., and Uematsu, S. (1992). High-frequency EEG activity at the start of seizures. *Journal of Clinical Neurophysiology*, 9(3):441–448.

Fitzpatrick, D., Lund, J. S., and Blasdel, G. G. (1985). Intrinsic connections of macaque striate cortex: afferent and efferent connections of lamina 4C. *Journal of Neuroscience Methods*, 5(12):3329–3349. PMID: 3001242.

Fountas, K. N. and Smith, J. R. (2007). A novel closed-loop stimulation system in the control of focal, medically refractory epilepsy. In Sakas, D. E. and Simpson, B. A., editors, *Operative Neuromodulation*, number 97/2 in Acta Neurochirurgica Supplements, pages 357–362. Springer Vienna.

Freund, T. F. and Katona, I. (2007). Perisomatic inhibition. *Neuron*, 56(1):33–42.

Frey, U., Egert, U., Heer, F., Hafizovic, S., and Hierlemann, A. (2009). Microelectronic system for high-resolution mapping of extracellular electric fields applied to brain slices. *Biosensors and Bioelectronics*, 24(7):2191–2198.

Fries, P. (2009). Neuronal gamma-band synchronization as a fundamental process in cortical computation. *Annual Review of Neuroscience*, 32(1):209–224.

Fries, P., Reynolds, J. H., Rorie, A. E., and Desimone, R. (2001). Modulation of oscillatory neuronal synchronization by selective visual attention. *Science*, 291(5508):1560–1563. PMID: 11222864.

Frostig, R. D. and Chen-Bee, C. H. (2012). Intrinsic signal optical imaging. In Brette, R. and Destexhe, A., editors, *Handbook of Neural Activity Measurement*, pages 92–135. Cambridge University Press, New York, NY, USA.

Fuchs, E. C., Zivkovic, A. R., Cunningham, M. O., Middleton, S., LeBeau, F. E. N., Bannerman, D. M., Rozov, A., Whittington, M. A., Traub, R. D., Rawlins, J. N. P., and Monyer, H. (2007). Recruitment of parvalbumin-positive interneurons determines hippocampal function and associated behavior. *Neuron*, 53(4):591–604.

Gabriel, S., Lau, R. W., and Gabriel, C. (1996). The dielectric properties of biological tissues: II. measurements in the frequency range 10 hz to 20 GHz. *Physics in Medicine and Biology*, 41(11):2251–2269. PMID: 8938025.

Gallinat, J., Winterer, G., Herrmann, C., and Senkowski, D. (2004). Reduced oscillatory gamma-band responses in unmedicated schizophrenic patients indicate impaired frontal network processing. *Clinical Neurophysiology*, 115(8):1863–1874.

Geisler, C., Brunel, N., and Wang, X.-J. (2005). Contributions of intrinsic membrane dynamics to fast network oscillations with irregular neuronal discharges. *Journal of neurophysiology*, 94(6):4344–4361. PMID: 16093332.

Gerstner, W. and Brette, R. (2009). Adaptive exponential integrate-and-fire model. *Scholarpedia*, 4(6):8427.

Gerstner, W. and Kistler, W. M. (2002). *Spiking Neuron Models: Single Neurons, Populations, Plasticity*. Cambridge University Press.

Gerstner, W., van Hemmen, J. L., and Cowan, J. D. (1996). What matters in neuronal locking? *Neural Computation*, 8(8):1653–1676.

Gewaltig, M.-O. and Diesmann, M. (2007). NEST (NEural simulation tool). *Scholarpedia*, 2(4):1430.

Gillespie, D. T. (1996). Exact numerical simulation of the ornstein-uhlenbeck process and its integral. *Physical Review E*, 54(2):2084.

Gleeson, P., Crook, S., Cannon, R. C., Hines, M. L., Billings, G. O., Farinella, M., Morse, T. M., Davison, A. P., Ray, S., Bhalla, U. S., Barnes, S. R., Dimitrova, Y. D., and Silver, R. A. (2010). NeuroML: a language for describing data driven models of neurons and networks with a high degree of biological detail. *PLOS Computational Biology*, 6(6):e1000815.

Glickfeld, L. L., Roberts, J. D., Somogyi, P., and Scanziani, M. (2009). Interneurons hyperpolarize pyramidal cells along their entire somatodendritic axis. *Nature Neuroscience*, 12(1):21–23.

Gold, C., Henze, D. A., and Koch, C. (2007). Using extracellular action potential recordings to constrain compartmental models. *Journal of Computational Neuroscience*, 23(1):39–58.

Gold, C., Henze, D. A., Koch, C., and Buzsáki, G. (2006). On the origin of the extracellular action potential waveform: A modeling study. *Journal of Neurophysiology*, 95(5):3113–3128. PMID: 16467426.

Goldenfeld, N. and Kadanoff, L. P. (1999). Simple lessons from complexity. *Science*, 284(5411):87–89. PMID: 10102823.

Goodman, D. F. M. and Brette, R. (2009). The brian simulator. *Frontiers in Neuroscience*, 3.

Gratiy, S. L., Devor, A., Einevoll, G. T., and Dale, A. M. (2011). On the estimation of population-specific synaptic currents from laminar multielectrode recordings. *Frontiers in Neuroinformatics*, 5:32. PMID: 22203801.

Gratiy, S. L., Pettersen, K. H., Einevoll, G. T., and Dale, A. M. (2013). Pitfalls in the interpretation of multielectrode data: on the infeasibility of the neuronal current-source monopoles. *Journal of Neurophysiology*, 109(6):1681–1682. PMID: 23503556.

Green, M., Mintz, J., Salveson, D., Nuechterlein, K., Breitmeyer, B., Light, G., and Braff, D. (2003). Visual masking as a probe for abnormal gamma range activity in schizophrenia. *Biological Psychiatry*, 53(12):1113–1119.

Grice, S., Spratling, M., Karmiloff-Smith, A., Halit, H., Csibra, G., De Haan, M., and Johnson, M. (2001). Disordered visual processing and oscillatory brain activity in autism and williams syndrome. *NeuroReport*, 12(12):2697–2700.

Grinvald, A. and Hildesheim, R. (2004). VSDI: a new era in functional imaging of cortical dynamics. *Nature Reviews Neuroscience*, 5(11):874–885.

Grinvald, A., Lieke, E., Frostig, R. D., Gilbert, C. D., and Wiesel, T. N. (1986). Functional architecture of cortex revealed by optical imaging of intrinsic signals. *Nature*, 324(6095):361–364.

Gruber, T., Tsivilis, D., Montaldi, D., and Müller, M. (2004). Induced gamma band responses: An early marker of memory encoding and retrieval. *NeuroReport*, 15(11):1837–1841.

Haas, L. (2003). Hans berger (1873-1941), richard caton (1842-1926), and electroencephalography. *Journal of Neurology, Neurosurgery, and Psychiatry*, 74(1):9. PMID: 12486257 PMCID: PMC1738204.

Haenschel, C., Baldeweg, T., Croft, R. J., Whittington, M., and Gruzelier, J. (2000). Gamma and beta frequency oscillations in response to novel auditory stimuli: A comparison of human electroencephalogram (EEG) data with in vitro models.

Proceedings of the National Academy of Sciences of the United States of America, 97(13):7645–7650. PMID: 10852953.

Haig, A., Gordon, E., De Pascalis, V., Meares, R., Bahramali, H., and Harris, A. (2000). Gamma activity in schizophrenia: evidence of impaired network binding? *Clinical Neurophysiology*, 111(8):1461–1468.

Hájos, N., Pálhalmi, J., Mann, E. O., Németh, B., Paulsen, O., and Freund, T. F. (2004). Spike timing of distinct types of GABAergic interneuron during hippocampal gamma oscillations in vitro. *The Journal of Neuroscience*, 24(41):9127–9137. PMID: 15483131.

Hájos, N. and Paulsen, O. (2009). Network mechanisms of gamma oscillations in the CA3 region of the hippocampus. *Neural Networks*, 22(8):1113–1119. PMID: 19683412.

Hämäläinen, M., Hari, R., Ilmoniemi, R. J., Knuutila, J., and Lounasmaa, O. V. (1993). Magnetoencephalography - theory, instrumentation, and applications to noninvasive studies of the working human brain. *Reviews of modern Physics*, 65(2):413.

Hashimoto, T., Volk, D., Eggan, S., Mirnics, K., Pierri, J., Sun, Z., Sampson, A., and Lewis, D. (2003). Gene expression deficits in a subclass of GABA neurons in the prefrontal cortex of subjects with schizophrenia. *The Journal of Neuroscience*, 23(15):6315–6326.

Hay, E., Hill, S., Schürmann, F., Markram, H., and Segev, I. (2011). Models of neocortical layer 5b pyramidal cells capturing a wide range of dendritic and perisomatic active properties. *PLOS Computational Biology*, 7(7):e1002107.

Hellwig, B. (2000). A quantitative analysis of the local connectivity between pyramidal neurons in layers 2/3 of the rat visual cortex. *Biological Cybernetics*, 82(2):111–121.

Henker, S., Partzsch, J., and Schüffny, R. (2011). Accuracy evaluation of numerical methods used in state-of-the-art simulators for spiking neural networks. *Journal of Computational Neuroscience*, 32(2):309–326.

Herrmann, C., Munk, M., and Engel, A. (2004). Cognitive functions of gamma-band activity: Memory match and utilization. *Trends in Cognitive Sciences*, 8(8):347–355.

Herrmann, C. S. and Demiralp, T. (2005). Human EEG gamma oscillations in neuropsychiatric disorders. *Clinical Neurophysiology*, 116(12):2719–2733.

Herz, A. V. M., Gollisch, T., Machens, C. K., and Jaeger, D. (2006). Modeling single-neuron dynamics and computations: A balance of detail and abstraction. *Science*, 314(5796):80–85.

Hille, B. (2001). *Ion channels of excitable membranes*, volume 507. Sinauer Sunderland, MA.

Hines, M. L. and Carnevale, N. T. (1997). The NEURON simulation environment. *Neural Computation*, 9(6):1179–1209. PMID: 9248061.

Hodgkin, A. L. and Huxley, A. F. (1952). A quantitative description of membrane current and its application to conduction and excitation in nerve. *The Journal of Physiology*, 117(4):500–544. PMID: 12991237 PMCID: PMC1392413.

Holt, G. R. (1998). *A critical reexamination of some assumptions and implications of cable theory in neurobiology*. PhD thesis, California Institute of Technology.

Holt, G. R. and Koch, C. (1999). Electrical interactions via the extracellular potential near cell bodies. *Journal of Computational Neuroscience*, 6(2):169–184. PMID: 10333161.

Hormuzdi, S. G., Pais, I., LeBeau, F. E. N., Towers, S. K., Rozov, A., Buhl, E. H., Whittington, M. A., and Monyer, H. (2001). Impaired electrical signaling disrupts gamma frequency oscillations in connexin 36-deficient mice. *Neuron*, 31(3):487–495.

House, P. A., MacDonald, J. D., Tresco, P. A., and Normann, R. A. (2006). Acute microelectrode array implantation into human neocortex: preliminary technique and histological considerations. *Neurosurgical Focus*, 20(5):1–4.

Izhikevich, E. (2004). Which model to use for cortical spiking neurons? *IEEE Transactions on Neural Networks*, 15(5):1063–1070.

Izhikevich, E. M. and Edelman, G. M. (2008). Large-scale model of mammalian thalamocortical systems. *Proceedings of the National Academy of Sciences*, 105(9):3593–3598. PMID: 18292226.

Jagust, W. J., Seab, J. P., Huesman, R. H., Valk, P. E., Mathis, C. A., Reed, B. R., Coxson, P. G., and Budinger, T. F. (1991). Diminished glucose transport in alzheimer's disease: Dynamic PET studies. *Journal of Cerebral Blood Flow and Metabolism*, 11(2):323–330.

Johnston, D. and Wu, S. M. S. (1995). *Foundations of Cellular Neurophysiology*. Massachusetts Institute of Technology Press.

Jonas, P. and Buzsaki, G. (2007). Neural inhibition. *Scholarpedia*, 2(9):3286.

Jutras, M. J., Fries, P., and Buffalo, E. A. (2009). Gamma-band synchronization in the macaque hippocampus and memory formation. *The Journal of Neuroscience*, 29(40):12521–12531. PMID: 19812327.

Kajikawa, Y. and Schroeder, C. E. (2011). How local is the local field potential? *Neuron*, 72(5):847–858.

Kandel, E., Schwartz, J., Jessell, T., Siegelbaum, S., and Hudspeth, A. J. (2012). *Principles of Neural Science, Fifth Edition*. McGraw Hill Professional.

Katzner, S., Nauhaus, I., Benucci, A., Bonin, V., Ringach, D. L., and Carandini, M. (2009). Local origin of field potentials in visual cortex. *Neuron*, 61(1):35–41. PMID: 19146811.

Kawaguchi, Y. and Kubota, Y. (1997). GABAergic cell subtypes and their synaptic connections in rat frontal cortex. *Cerebral Cortex*, 7(6):476–486. PMID: 9276173.

Kawaguchi, Y. and Kubota, Y. (1998). Neurochemical features and synaptic connections of large physiologically-identified GABAergic cells in the rat frontal cortex. *Neuroscience*, 85(3):677–701.

Kissler, J., Müller, M., Fehr, T., Rockstroh, B., and Elbert, T. (2000). MEG gamma band activity in schizophrenia patients and healthy subjects in a mental arithmetic task and at rest. *Clinical Neurophysiology*, 111(11):2079–2087.

Kisvárday, Z. F. and Eysel, U. T. (1992). Cellular organization of reciprocal patchy networks in layer III of cat visual cortex (area 17). *Neuroscience*, 46(2):275–286. PMID: 1542406.

Koch, C. (1999). *Biophysics of computation: information processing in single neurons*. Oxford University Press, New York.

Koch, C. and Segev, I. (1998). *Methods in Neuronal Modeling: From Ions to Networks*. Massachusetts Institute of Technology Press.

Koch, S. P., Werner, P., Steinbrink, J., Fries, P., and Obrig, H. (2009). Stimulus-induced and state-dependent sustained gamma activity is tightly coupled to the hemodynamic response in humans. *The Journal of Neuroscience*, 29(44):13962–13970. PMID: 19890006.

Koenig, T., Prichep, L., Dierks, T., Hubl, D., Wahlund, L., John, E., and Jelic, V. (2005). Decreased EEG synchronization in alzheimer's disease and mild cognitive impairment. *Neurobiology of Aging*, 26(2):165–171.

Kopell, N. and Ermentrout, G. B. (1988). Coupled oscillators and the design of central pattern generators. *Mathematical Biosciences*, 90(1–2):87–109.

Kopell, N. and Ermentrout, G. B. (2002). Chapter 1 mechanisms of phase-locking and frequency control in pairs of coupled neural oscillators. In Bernold Fiedler, editor, *Handbook of Dynamical Systems*, volume Volume 2, pages 3–54. Elsevier Science.

Kopell, N., Ermentrout, G. B., Whittington, M. A., and Traub, R. D. (2000). Gamma rhythms and beta rhythms have different synchronization properties. *Proceedings of the National Academy of Sciences of the United States of America*, 97(4):1867–1872. PMID: 10677548 PMCID: PMC26528.

Kosaka, T., Katsumaru, H., Hama, K., Wu, J.-Y., and Heizmann, C. W. (1987). GABAergic neurons containing the ca2+-binding protein parvalbumin in the rat hippocampus and dentate gyrus. *Brain Research*, 419(1–2):119–130.

Kramer, M. A., Roopun, A. K., Carracedo, L. M., Traub, R. D., Whittington, M. A., and Kopell, N. J. (2008). Rhythm generation through period concatenation in

rat somatosensory cortex. *PLoS computational biology*, 4(9):e1000169. PMID: 18773075 PMCID: PMC2518953.

Kreiman, G., Hung, C. P., Kraskov, A., Quiroga, R. Q., Poggio, T., and DiCarlo, J. J. (2006). Object selectivity of local field potentials and spikes in the macaque inferior temporal cortex. *Neuron*, 49(3):433–445.

Kuhl, D. E., Koeppe, R. A., Minoshima, S., Snyder, S. E., Ficaro, E. P., Foster, N. L., Frey, K. A., and Kilbourn, M. R. (1999). In vivo mapping of cerebral acetylcholinesterase activity in aging and alzheimer's disease. *Neurology*, 52(4):691–699. PMID: 10078712.

Kuramoto, Y. (1975). Self-entrainment of a population of coupled non-linear oscillators. In Araki, P. H., editor, *International Symposium on Mathematical Problems in Theoretical Physics*, number 39 in Lecture Notes in Physics, pages 420–422. Springer Berlin Heidelberg.

LaLumiere, R. T. (2011). A new technique for controlling the brain: optogenetics and its potential for use in research and the clinic. *Brain Stimulation*, 4(1):1–6. PMID: 21255749.

Lee, K.-H., Williams, L., Breakspear, M., and Gordon, E. (2003a). Synchronous gamma activity: A review and contribution to an integrative neuroscience model of schizophrenia. *Brain Research Reviews*, 41(1):57–78.

Lee, K.-H., Williams, L., Haig, A., and Gordon, E. (2003b). "Gamma (40 hz) phase synchronicity" and symptom dimensions in schizophrenia. *Cognitive Neuropsychiatry*, 8(1):57–71. PMID: 16571550.

Lefort, S., Tomm, C., Floyd Sarria, J.-C., and Petersen, C. C. (2009). The excitatory neuronal network of the c2 barrel column in mouse primary somatosensory cortex. *Neuron*, 61(2):301–316.

Leiderman, D. B., Balish, M., Sato, S., Kufta, C., Reeves, P., Gaillard, W. D., and Theodore, W. H. (1992). Comparison of PET measurements of cerebral blood flow and glucose metabolism for the localization of human epileptic foci. *Epilepsy Research*, 13(2):153–157.

Łęski, S., Lindén, H., Tetzlaff, T., Pettersen, K. H., and Einevoll, G. T. (2013). Frequency dependence of signal power and spatial reach of the local field potential. *PLOS Computational Biology*, 9(7):e1003137. PMID: 23874180.

Łęski, S., Pettersen, K. H., Tunstall, B., Einevoll, G. T., Gigg, J., and Wójcik, D. K. (2011). Inverse current source density method in two dimensions: inferring neural activation from multielectrode recordings. *Neuroinformatics*, 9(4):401–425. PMID: 21409556.

Lewis, D. A., Curley, A. A., Glausier, J. R., and Volk, D. W. (2012). Cortical parvalbumin interneurons and cognitive dysfunction in schizophrenia. *Trends in Neurosciences*, 35(1):57–67.

Light, G. A., Hsu, J. L., Hsieh, M. H., Meyer-Gomes, K., Sprock, J., Swerdlow, N. R., and Braff, D. L. (2006). Gamma band oscillations reveal neural network cortical coherence dysfunction in schizophrenia patients. *Biological Psychiatry*, 60(11):1231–1240.

Lindén, H., Hagen, E., Łęski, S., Norheim, E. S., Pettersen, K. H., and Einevoll, G. T. (2014). LFPy: a tool for biophysical simulation of extracellular potentials generated by detailed model neurons. *Frontiers in Neuroinformatics*, 7:41.

Lindén, H., Pettersen, K. H., and Einevoll, G. T. (2010). Intrinsic dendritic filtering gives low-pass power spectra of local field potentials. *Journal of Computational Neuroscience*, 29(3):423–444.

Lindén, H., Tetzlaff, T., Potjans, T. C., Pettersen, K. H., Grün, S., Diesmann, M., and Einevoll, G. T. (2011). Modeling the spatial reach of the LFP. *Neuron*, 72(5):859–872.

Liu, J. and Newsome, W. T. (2006). Local field potential in cortical area MT: stimulus tuning and behavioral correlations. *The Journal of Neuroscience*, 26(30):7779–7790. PMID: 16870724.

Logothetis, N. K. (2003). The underpinnings of the BOLD functional magnetic resonance imaging signal. *The Journal of Neuroscience*, 23(10):3963–3971.

Logothetis, N. K. (2008). What we can do and what we cannot do with fMRI. *Nature*, 453(7197):869–878.

Logothetis, N. K., Kayser, C., and Oeltermann, A. (2007). In vivo measurement of cortical impedance spectrum in monkeys: Implications for signal propagation. *Neuron*, 55(5):809–823.

Logothetis, N. K., Pauls, J., Augath, M., Trinath, T., and Oeltermann, A. (2001). Neurophysiological investigation of the basis of the fMRI signal. *Nature*, 412(6843):150–157.

Logothetis, N. K. and Wandell, B. A. (2004). Interpreting the BOLD signal. *Annual Review of Physiology*, 66:735–769. PMID: 14977420.

Lopes da Silva, F. (2013). EEG and MEG: relevance to neuroscience. *Neuron*, 80(5):1112–1128. PMID: 24314724.

Mainen, Z. F. and Sejnowski, T. J. (1996). Influence of dendritic structure on firing pattern in model neocortical neurons. *Nature*, 382(6589):363–366. PMID: 8684467.

Mainy, N., Kahane, P., Minotti, L., Hoffmann, D., Bertrand, O., and Lachaux, J.-P. (2007). Neural correlates of consolidation in working memory. *Human Brain Mapping*, 28(3):183–193.

Mann, E. O., Suckling, J. M., Hajos, N., Greenfield, S. A., and Paulsen, O. (2005). Perisomatic feedback inhibition underlies cholinergically induced fast network oscillations in the rat hippocampus in vitro. *Neuron*, 45(1):105–117.

Marco, P. and DeFelipe, J. (1997). Altered synaptic circuitry in the human temporal neocortex removed from epileptic patients. *Experimental Brain Research*, 114(1):1–10.

Marco, P., Sola, R. G., Ramón y Cajal, S., and DeFelipe, J. (1997). Loss of inhibitory synapses on the soma and axon initial segment of pyramidal cells in human epileptic peritumoural neocortex: implications for epilepsy. *Brain Research Bulletin*, 44(1):47–66. PMID: 9288831.

Markram, H. (2006). The blue brain project. *Nature reviews. Neuroscience*, 7(2):153–160. PMID: 16429124.

Markram, H., Toledo-Rodriguez, M., Wang, Y., Gupta, A., Silberberg, G., and Wu, C. (2004). Interneurons of the neocortical inhibitory system. *Nature Reviews Neuroscience*, 5(10):793–807.

Mazzoni, A., Panzeri, S., Logothetis, N. K., and Brunel, N. (2008). Encoding of naturalistic stimuli by local field potential spectra in networks of excitatory and inhibitory neurons. *PLOS Computational Biology*, 4(12):e1000239.

Mazzoni, A., Whittingstall, K., Brunel, N., Logothetis, N. K., and Panzeri, S. (2010). Understanding the relationships between spike rate and delta/gamma frequency bands of LFPs and EEGs using a local cortical network model. *NeuroImage*, 52(3):956–972.

Menendez de la Prida, L. and Trevelyan, A. J. (2011). Cellular mechanisms of high frequency oscillations in epilepsy: On the diverse sources of pathological activities. *Epilepsy Research*, 97(3):308–317.

Mitra, P. and Bokil, H. (2008). *Observed brain dynamics*. Oxford University Press, Oxford; New York.

Mitzdorf, U. (1985). Current source-density method and application in cat cerebral cortex: investigation of evoked potentials and EEG phenomena. *Physiological Reviews*, 65(1):37–100. PMID: 3880898.

Molnár, M., Gács, G., Újvári, G., Skinner, J., and Karmos, G. (1997). Dimensional complexity of the EEG in subcortical stroke - a case study. *International Journal of Psychophysiology*, 25(3):193–199.

Morrison, A., Mehring, C., Geisel, T., Aertsen, T. G. A., and Diesmann, M. A. (2005). Advancing the boundaries of high-connectivity network simulation with distributed computing. *Neural Computation*, 17(8):1776–1801.

Mosher, J., Leahy, R., and Lewis, P. (1999). EEG and MEG: forward solutions for inverse methods. *IEEE Transactions on Biomedical Engineering*, 46(3):245–259.

Naud, R., Marcille, N., Clopath, C., and Gerstner, W. (2008). Firing patterns in the adaptive exponential integrate-and-fire model. *Biological Cybernetics*, 99(4-5):335–347.

Nicholson, C. and Freeman, J. A. (1975). Theory of current source-density analysis and determination of conductivity tensor for anuran cerebellum. *Journal of Neurophysiology*, 38(2):356–368. PMID: 805215.

Niebur, E. (2008). Neuronal cable theory. *Scholarpedia*, 3(5):2674.

Niedermeyer, E. and Silva, F. H. L. d. (2005). *Electroencephalography: Basic Principles, Clinical Applications, and Related Fields*. Lippincott Williams & Wilkins.

Niessing, J., Ebisch, B., Schmidt, K. E., Niessing, M., Singer, W., and Galuske, R. A. W. (2005). Hemodynamic signals correlate tightly with synchronized gamma oscillations. *Science*, 309(5736):948–951. PMID: 16081740.

Nunez, P. L. and Srinivasan, R. (2006). *Electric Fields of the Brain: The Neurophysics of EEG*. Oxford University Press, New York, NY, USA, 2nd edition edition.

Oren, I., Hájos, N., and Paulsen, O. (2010). Identification of the current generator underlying cholinergically induced gamma frequency field potential oscillations in the hippocampal CA3 region. *The Journal of Physiology*, 588(5):785–797. PMID: 20051494.

Oren, I., Mann, E. O., Paulsen, O., and Hajos, N. (2006). Synaptic currents in anatomically identified CA3 neurons during hippocampal gamma oscillations in vitro. *The Journal of Neuroscience*, 26(39):9923–9934. PMID: 17005856.

Osipova, D., Takashima, A., Oostenveld, R., Fernández, G., Maris, E., and Jensen, O. (2006). Theta and gamma oscillations predict encoding and retrieval of declarative memory. *The Journal of Neuroscience*, 26(28):7523–7531. PMID: 16837600.

Pedreira, C., Martinez, J., Ison, M. J., and Quian Quiroga, R. (2012). How many neurons can we see with current spike sorting algorithms? *Journal of Neuroscience Methods*, 211(1):58–65.

Percival, D. B. and Walden, A. T. (1993). *Spectral Analysis for Physical Applications*. Cambridge University Press.

Peters, A. and Payne, B. R. (1993). Numerical relationships between geniculocortical afferents and pyramidal cell modules in cat primary visual cortex. *Cerebral Cortex*, 3(1):69–78. PMID: 8439740.

Petilla Interneuron Nomenclature Group, Ascoli, G. A., Alonso-Nanclares, L., Anderson, S. A., Barrionuevo, G., Benavides-Piccione, R., Burkhalter, A., Buzsáki, G., Cauli, B., Defelipe, J., Fairén, A., Feldmeyer, D., Fishell, G., Fregnac, Y., Freund, T. F., Gardner, D., Gardner, E. P., Goldberg, J. H., Helmstaedter, M., Hestrin, S., Karube, F., Kisvárday, Z. F., Lambolez, B., Lewis, D. A., Marin, O., Markram, H., Muñoz, A., Packer, A., Petersen, C. C. H., Rockland, K. S., Rossier, J., Rudy, B., Somogyi, P., Staiger, J. F., Tamas, G., Thomson, A. M., Toledo-Rodriguez, M., Wang, Y., West, D. C., and Yuste, R. (2008). Petilla terminology: nomenclature of features of GABAergic interneurons of the cerebral cortex. *Nature reviews. Neuroscience*, 9(7):557–568. PMID: 18568015 PMCID: PMC2868386.

Pettersen, K. H., Devor, A., Ulbert, I., Dale, A. M., and Einevoll, G. T. (2006). Current-source density estimation based on inversion of electrostatic forward solution: Effects of finite extent of neuronal activity and conductivity discontinuities. *Journal of Neuroscience Methods*, 154(1–2):116–133.

Pettersen, K. H. and Einevoll, G. T. (2008). Amplitude variability and extracellular low-pass filtering of neuronal spikes. *Biophysical Journal*, 94(3):784–802.

Pettersen, K. H., Hagen, E., and Einevoll, G. T. (2008). Estimation of population firing rates and current source densities from laminar electrode recordings. *Journal of Computational Neuroscience*, 24(3):291–313.

Pettersen, K. H., Lindén, H., Dale, A. M., and Einevoll, G. T. (2012). Extracellular spikes and current-source density. In Brette, R. and Destexhe, A., editors, *Handbook of Neural Activity Measurement*, pages 92–135. Cambridge University Press, New York, NY, USA.

Plonsey, R. and Barr, R. C. (2000). *Biolelectricity: A Quantitative Approach*. Springer, New York, NY, USA, 2nd edition edition.

Plonsey, R. and Heppner, D. B. (1967). Considerations of quasi-stationarity in electrophysiological systems. *The bulletin of mathematical biophysics*, 29(4):657–664.

Potjans, T. C. and Diesmann, M. (2012). The cell-type specific cortical microcircuit:

Relating structure and activity in a full-scale spiking network model. *Cerebral Cortex*.

Powell, S. B., Sejnowski, T. J., and Behrens, M. M. (2012). Behavioral and neurochemical consequences of cortical oxidative stress on parvalbumin-interneuron maturation in rodent models of schizophrenia. *Neuropharmacology*, 62(3):1322–1331.

Press, W. H., Teukolsky, S. A., Vetterling, W. T., and Flannery, B. P. (2007). *Numerical Recipes 3rd Edition: The Art of Scientific Computing*. Cambridge University Press.

Prinz, A. A., Bucher, D., and Marder, E. (2004). Similar network activity from disparate circuit parameters. *Nature neuroscience*, 7(12):1345–1352. PMID: 15558066.

Quian Quiroga, R. (2007). Spike sorting. *Scholarpedia*, 2(12):3583.

Quian Quiroga, R., Nadasdy, Z., and Ben-Shaul, Y. (2004). Unsupervised spike detection and sorting with wavelets and superparamagnetic clustering. *Neural Computation*, 16(8):1661–1687.

Raichle, M. E. and Mintun, M. A. (2006). Brain work and brain imaging. *Annual Review of Neuroscience*, 29(1):449–476. PMID: 16776593.

Ray, S. and Maunsell, J. (2010). Differences in gamma frequencies across visual cortex restrict their possible use in computation. *Neuron*, 67(5):885–896.

Reimann, M. W., Anastassiou, C. A., Perin, R., Hill, S. L., Markram, H., and Koch, C. (2013). A biophysically detailed model of neocortical local field potentials predicts the critical role of active membrane currents. *Neuron*, 79(2):375–390. PMID: 23889937.

Rice, J. K., Rorden, C., Little, J. S., and Parra, L. C. (2013). Subject position affects EEG magnitudes. *NeuroImage*, 64:476–484.

Riera, J. J., Ogawa, T., Goto, T., Sumiyoshi, A., Nonaka, H., Evans, A., Miyakawa, H., and Kawashima, R. (2012). Pitfalls in the dipolar model for the neocortical EEG sources. *Journal of Neurophysiology*, 108(4):956–975.

Roopun, A. K., Kramer, M. A., Carracedo, L. M., Kaiser, M., Davies, C. H., Traub, R. D., Kopell, N. J., and Whittington, M. A. (2008). Period concatenation underlies interactions between gamma and beta rhythms in neocortex. *Frontiers in cellular neuroscience*, 2:1. PMID: 18946516 PMCID: PMC2525927.

Roopun, A. K., Simonotto, J. D., Pierce, M. L., Jenkins, A., Nicholson, C., Schofield, I. S., Whittaker, R. G., Kaiser, M., Whittington, M. A., Traub, R. D., and Cunningham, M. O. (2010). A nonsynaptic mechanism underlying interictal discharges in human epileptic neocortex. *Proceedings of the National Academy of Sciences of the United States of America*, 107(1):338–343. PMID: 19966298 PMCID: PMC2806783.

Rosenfalck, P. (1969). Intra- and extracellular potential fields of active nerve and muscle fibres. a physico-mathematical analysis of different models. *Thrombosis et diathesis haemorrhagica. Supplementum*, 321:1–168. PMID: 5383732.

Roxin, A. (2011). The role of degree distribution in shaping the dynamics in networks of sparsely connected spiking neurons. *Frontiers in Computational Neuroscience*, 5.

Roxin, A., Brunel, N., and Hansel, D. (2005). Role of delays in shaping spatiotemporal dynamics of neuronal activity in large networks. *Physical Review Letters*, 94(23):238103.

Schevon, C. A., Ng, S. K., Cappell, J., Goodman, R. R., McKhann, G., Waziri, A., Branner, A., Sosunov, A., Schroeder, C. E., and Emerson, R. G. (2008). Microphysiology of epileptiform activity in human neocortex. *Journal of Clinical Neurophysiology*, 25(6):321–330.

Schevon, C. A., Trevelyan, A. J., Schroeder, C. E., Goodman, R. R., McKhann, G., and Emerson, R. G. (2009). Spatial characterization of interictal high frequency oscillations in epileptic neocortex. *Brain*, 132(11):3047–3059.

Schomburg, E. W., Anastassiou, C. A., Buzsáki, G., and Koch, C. (2012). The spiking component of oscillatory extracellular potentials in the rat hippocampus. *The Journal of Neuroscience*, 32(34):11798–11811. PMID: 22915121.

Scotti, A., Kalt, G., Bollag, O., and Nitsch, C. (1997). Parvalbumin disappears from GABAergic CA1 neurons of the gerbil hippocampus with seizure onset while its presence persists in the perforant path. *Brain Research*, 760(1-2):109–117.

Sederberg, P., Schulze-Bonhage, A., Madsen, J., Bromfield, E., McCarthy, D., Brandt, A., Tully, M., and Kahana, M. (2007). Hippocampal and neocortical gamma oscillations predict memory formation in humans. *Cerebral Cortex*, 17(5):1190–1196.

Shepherd, G. M., editor (2003). *The Synaptic Organization of the Brain*. Oxford University Press.

Simon, A., Traub, R. D., Vladimirov, N., Jenkins, A., Nicholson, C., Whittaker, R. G., Schofield, I., Clowry, G. J., Cunningham, M. O., and Whittington, M. A. (2014). Gap junction networks can generate both ripple-like and fast ripple-like oscillations. *European Journal of Neuroscience*, 39(1):46–60.

Slewa-Younan, S., Green, A., Baguley, I., Felmingham, K., Haig, A., and Gordon, E. (2002). Is 'gamma' (40 hz) synchronous activity disturbed in patients with traumatic brain injury? *Clinical Neurophysiology*, 113(10):1640–1646.

Smith, J. R., Fountas, K. N., Murro, A. M., Park, Y. D., Jenkins, P. D., Morrell, M., Esteller, R., and Greene, D. (2010). Closed-loop stimulation in the control of focal epilepsy of insular origin. *Stereotactic and Functional Neurosurgery*, 88(5):281–287. PMID: 20588079.

Spencer, K., Nestor, P., Perlmutter, R., Niznikiewicz, M., Klump, M., Frumin, M., Shenton, M., and McCarley, R. (2004). Neural synchrony indexes disordered perception and cognition in schizophrenia. *Proceedings of the National Academy of Sciences of the United States of America*, 101(49):17288–17293.

Spencer, K. M. (2009a). The functional consequences of cortical circuit abnormalities on gamma oscillations in schizophrenia: Insights from computational modeling. *Frontiers in Human Neuroscience*, 3. PMID: 19876408 PMCID: PMC2769552.

Spencer, K. M. (2009b). The functional consequences of cortical circuit abnormalities on gamma oscillations in schizophrenia: Insights from computational modeling. *Frontiers in Human Neuroscience*, 3. PMID: 19876408 PMCID: PMC2769552.

Staley, K., Swiercz, W., and Cios, K. J. (2011). Validating models of epilepsy. In Soltesz, I. and Staley, K., editors, *Computational Neuroscience in Epilepsy*. Academic Press.

Stam, C., Van Cappellen van Walsum, A., Pijnenburg, Y., Berendse, H., De Munck, J., Scheltens, P., and Van Dijk, B. (2002). Generalized synchronization of MEG recordings in alzheimer's disease: Evidence for involvement of the gamma band. *Journal of Clinical Neurophysiology*, 19(6):562–574.

Stepanyants, A., Hirsch, J. A., Martinez, L. M., Kisvarday, Z. F., Ferecsko, A. S., and Chklovskii, D. B. (2008). Local potential connectivity in cat primary visual cortex. *Cerebral Cortex*, 18(1):13–28.

Stevenson, I. H. and Kording, K. P. (2011). How advances in neural recording affect data analysis. *Nature Neuroscience*, 14(2):139–142.

Strong, A. J., Fabricius, M., Boutelle, M. G., Hibbins, S. J., Hopwood, S. E., Jones, R., Parkin, M. C., and Lauritzen, M. (2002). Spreading and synchronous depressions of cortical activity in acutely injured human brain. *Stroke*, 33(12):2738–2743. PMID: 12468763.

Sun, Y., Farzan, F., Barr, M. S., Kirihara, K., Fitzgerald, P. B., Light, G. A., and Daskalakis, Z. J. (2011). Gamma oscillations in schizophrenia: Mechanisms and clinical significance. *Brain Research*, 1413:98–114.

Tallon-Baudry, C., Bertrand, O., Hénaff, M.-A., Isnard, J., and Fischer, C. (2005). Attention modulates gamma-band oscillations differently in the human lateral occipital cortex and fusiform gyrus. *Cerebral Cortex*, 15(5):654–662.

Tallon-Baudry, C., Bertrand, O., Peronnet, F., and Pernier, J. (1998). Induced γ-band activity during the delay of a visual short-term memory task in humans. *The Journal of Neuroscience*, 18(11):4244–4254.

Tam, A. T. C. and Wang, C.-L. (2000). Efficient scheduling of complete exchange on clusters. In Chaudhry, G. and Sha, E., editors, *13th International Conference on Parallel and Distributed Computing Systems (PDCS 2000)*, pages 111–116, Las Vegas, Nevada. International Society for Computers and Their Applications.

Tekell, J. L., Hoffmann, R., Hendrickse, W., Greene, R. W., Rush, A. J., and Armitage, R. (2005). High frequency EEG activity during sleep: Characteristics in schizophrenia and depression. *Clinical EEG and Neuroscience*, 36(1):25–35. PMID: 15683195.

Thomson, A. M. and Lamy, C. (2007). Functional maps of neocortical local circuitry. *Frontiers in Neuroscience*, 1(1):19.

Thomson, A. M., West, D. C., Wang, Y., and Bannister, A. P. (2002). Synaptic connections and small circuits involving excitatory and inhibitory neurons in layers 2-5 of adult rat and cat neocortex: Triple intracellular recordings and biocytin labelling in vitro. *Cerebral Cortex*, 12(9):936–953.

Thomson, D. (1982). Spectrum estimation and harmonic analysis. *Proceedings of the IEEE*, 70(9):1055–1096.

Thomson, W. (1854). On the theory of the electric telegraph. *Proceedings of the Royal Society of London*, 7:382–399. ArticleType: research-article / Full publication date: 1854 - 1855 /.

Touboul, J. and Brette, R. (2008). Dynamics and bifurcations of the adaptive exponential integrate-and-fire model. *Biological Cybernetics*, 99(4-5):319–334.

Traub, R. D., Bibbig, A., LeBeau, F. E. N., Cunningham, M. O., and Whittington, M. A. (2005a). Persistent gamma oscillations in superficial layers of rat auditory neocortex: experiment and model. *The Journal of Physiology*, 562(Pt 1):3–8. PMID: 15489250.

Traub, R. D., Buhl, E. H., Gloveli, T., and Whittington, M. A. (2003a). Fast rhythmic bursting can be induced in layer 2/3 cortical neurons by enhancing persistent na+ conductance or by blocking BK channels. *Journal of Neurophysiology*, 89(2):909–921. PMID: 12574468.

Traub, R. D., Contreras, D., Cunningham, M. O., Murray, H., LeBeau, F. E. N., Roopun, A., Bibbig, A., Wilent, W. B., Higley, M. J., and Whittington, M. A. (2005b). Single-column thalamocortical network model exhibiting gamma oscillations, sleep spindles, and epileptogenic bursts. *Journal of Neurophysiology*, 93(4):2194–2232. PMID: 15525801.

Traub, R. D., Cunningham, M. O., Gloveli, T., LeBeau, F. E. N., Bibbig, A., Buhl, E. H., and Whittington, M. A. (2003b). GABA-enhanced collective behavior in neuronal axons underlies persistent gamma-frequency oscillations. *Proceedings of the National Academy of Sciences of the United States of America*, 100(19):11047–11052. PMID: 12960382.

Traub, R. D., Jefferys, J. G. R., and Whittington, M. A. (1999a). *Fast Oscillations in Cortical Circuits*. MIT Press.

Traub, R. D., Kopell, N., Bibbig, A., Buhl, E. H., LeBeau, F. E. N., and Whittington, M. A. (2001a). Gap junctions between interneuron dendrites can enhance synchrony of gamma oscillations in distributed networks. *The Journal of Neuroscience*, 21(23):9478–9486. PMID: 11717382.

Traub, R. D., Miles, R., and Wong, R. K. (1989). Model of the origin of rhythmic population oscillations in the hippocampal slice. *Science (New York, N.Y.)*, 243(4896):1319–1325. PMID: 2646715.

Traub, R. D., Schmitz, D., Jefferys, J. G., and Draguhn, A. (1999b). High-frequency population oscillations are predicted to occur in hippocampal pyramidal neuronal networks interconnected by axoaxonal gap junctions. *Neuroscience*, 92(2):407–426. PMID: 10408594.

Traub, R. D. and Whittington, M. A. (2010). *Cortical Oscillations in Health and Disease*. Oxford University Press.

Traub, R. D., Whittington, M. A., Buhl, E. H., LeBeau, F. E. N., Bibbig, A., Boyd, S., Cross, H., and Baldeweg, T. (2001b). A possible role for gap junctions in generation of very fast EEG oscillations preceding the onset of, and perhaps initiating, seizures. *Epilepsia*, 42(2):153–170.

Traub, R. D., Whittington, M. A., Colling, S. B., Buzsáki, G., and Jefferys, J. G. (1996a). Analysis of gamma rhythms in the rat hippocampus in vitro and in vivo. *The Journal of Physiology*, 493 (Pt 2):471–484. PMID: 8782110.

Traub, R. D., Whittington, M. A., Stanford, I. M., and Jefferys, J. G. R. (1996b). A mechanism for generation of long-range synchronous fast oscillations in the cortex. *Nature*, 383(6601):621–624.

Trevelyan, A. J. (2009). The direct relationship between inhibitory currents and local field potentials. *The Journal of Neuroscience*, 29(48):15299–15307. PMID: 19955382.

van der Linden, S., Panzica, F., and Curtis, M. d. (1999). Carbachol induces fast oscillations in the medial but not in the lateral entorhinal cortex of the isolated guinea pig brain. *Journal of Neurophysiology*, 82(5):2441–2450. PMID: 10561417.

Voges, N., Guijarro, C., Aertsen, A., and Rotter, S. (2009). Models of cortical networks with long-range patchy projections. *Journal of Computational Neuroscience*, 28(1):137–154.

Voges, N., Schüz, A., Aertsen, A., and Rotter, S. (2010). A modeler's view on the spatial structure of intrinsic horizontal connectivity in the neocortex. *Progress in Neurobiology*, 92(3):277–292.

Volman, V., Behrens, M. M., and Sejnowski, T. J. (2011). Downregulation of parvalbumin at cortical GABA synapses reduces network gamma oscillatory activity. *The Journal of Neuroscience*, 31(49):18137–18148. PMID: 22159125.

Waldert, S., Lemon, R. N., and Kraskov, A. (2013). Influence of spiking activity on cortical local field potentials. *The Journal of Physiology*, 591(21):5291–5303. PMID: 23981719.

Wang, X.-J. (2010). Neurophysiological and computational principles of cortical rhythms in cognition. *Physiological Reviews*, 90(3):1195–1268.

Wang, X.-J. and Rinzel, J. (1992). Alternating and synchronous rhythms in reciprocally inhibitory model neurons. *Neural Computation*, 4(1):84–97.

Weckstrom, M. (2010). Intracellular recording. *Scholarpedia*, 5(8):2224.

White, J. A., Chow, C. C., Rit, J., Soto-Treviño, C., and Kopell, N. (1998). Synchronization and oscillatory dynamics in heterogeneous, mutually inhibited neurons. *Journal of Computational Neuroscience*, 5(1):5–16.

Whittington, M. A., Cunningham, M. O., LeBeau, F. E. N., Racca, C., and Traub, R. D. (2011). Multiple origins of the cortical gamma rhythm. *Developmental Neurobiology*, 71(1):92–106. PMID: 21154913.

Whittington, M. A., Doheny, H. C., Traub, R. D., LeBeau, F. E. N., and Buhl, E. H. (2001). Differential expression of synaptic and nonsynaptic mechanisms underlying stimulus-induced gamma oscillations in vitro. *The Journal of Neuroscience*, 21(5):1727–1738. PMID: 11222662.

Whittington, M. A., Stanford, I. M., Colling, S. B., Jefferys, J. G. R., and Traub, R. D. (1997a). Spatiotemporal patterns of γ frequency oscillations tetanically induced in the rat hippocampal slice. *The Journal of Physiology*, 502(3):591–607.

Whittington, M. A., Traub, R. D., Faulkner, H. J., Stanford, I. M., and Jefferys, J. G. R. (1997b). Recurrent excitatory postsynaptic potentials induced by synchronized fast cortical oscillations. *Proceedings of the National Academy of Sciences of the United States of America*, 94(22):12198–12203. PMID: 9342386.

Whittington, M. A., Traub, R. D., and Jefferys, J. G. R. (1995). Synchronized oscillations in interneuron networks driven by metabotropic glutamate receptor activation. *Nature*, 373(6515):612–615.

Whittington, M. A., Traub, R. D., Kopell, N., Ermentrout, B., and Buhl, E. H. (2000). Inhibition-based rhythms: experimental and mathematical observations on network dynamics. *International Journal of Psychophysiology*, 38(3):315–336.

Willoughby, J. O., Fitzgibbon, S. P., Pope, K. J., Mackenzie, L., Medvedev, A. V., Clark, C. R., Davey, M. P., and Wilcox, R. A. (2003). Persistent abnormality detected in the non-ictal electroencephalogram in primary generalised epilepsy. *Journal of Neurology, Neurosurgery & Psychiatry*, 74(1):51–55. PMID: 12486266.

Winfree, A. T. (1967). Biological rhythms and the behavior of populations of coupled oscillators. *Journal of Theoretical Biology*, 16(1):15–42.

Worrell, G. A., Parish, L., Cranstoun, S. D., Jonas, R., Baltuch, G., and Litt, B. (2004). High-frequency oscillations and seizure generation in neocortical epilepsy. *Brain*, 127(7):1496–1506. PMID: 15155522.

Xing, D., Yeh, C.-I., and Shapley, R. M. (2009). Spatial spread of the local field potential and its laminar variation in visual cortex. *The Journal of Neuroscience*, 29(37):11540–11549. PMID: 19759301.

Zaitsev, A. V., Gonzalez-Burgos, G., Povysheva, N. V., Kröner, S., Lewis, D. A., and Krimer, L. S. (2005). Localization of calcium-binding proteins in physiologically and morphologically characterized interneurons of monkey dorsolateral prefrontal cortex. *Cerebral Cortex*, 15(8):1178–1186. PMID: 15590911.

Zanos, T. P., Mineault, P. J., and Pack, C. C. (2011). Removal of spurious correlations between spikes and local field potentials. *Journal of Neurophysiology*, 105(1):474–486. PMID: 21068271.

Zhang, M., Ladas, T. P., Qiu, C., Shivacharan, R. S., Gonzalez-Reyes, L. E., and Durand, D. M. (2014). Propagation of epileptiform activity can be independent of synaptic transmission, gap junctions, or diffusion and is consistent with electrical field transmission. *The Journal of Neuroscience*, 34(4):1409–1419. PMID: 24453330.

Appendix A

Parameter tables for the cortical slice model

Extra parameters for the neocortical slice model described in Chapter 4.

Postsynaptic neurons		Presynaptic neurons: percent of cells	max no. synapses	p23	b23	nb23	ss4(L4)	ss4(L23)	p4	b4	nb4	p5(L23)	p5(L56)	b5	nb5	p6(L4)	p6(L56)	b6
P23	L23	26.3	5773	60.1	9.2	4.9	0.6	6.9	7.8	0.8	-	7.5	-	-	-	2.4	-	-
	L1		87	95.1	1.6	-	-	0.3	1.5	0.1	-	1.1	-	-	-	-	-	-
B23		3.1	3702	53.7	11.0	12.3	0.5	6.1	6.8	0.9	-	6.6	-	-	-	2.1	-	-
NB23		4.2	3144	57.8	12.6	4.6	0.5	6.6	7.5	0.9	-	7.2	-	-	-	2.2	-	-
SS4(L4)		9.3	4113	3.8	0.3	-	16.7	5.2	5.8	12.9	8.1	1.1	0.1	-	-	46.1	-	-
SS4(L23)		9.3	3610	7.7	0.6	-	15.6	5.3	5.9	12.8	7.6	1.5	0.1	-	-	43.0	-	-
P4	L4	9.3	3619	6.0	0.3	-	16.0	5.0	5.8	12.9	8.3	1.6	0.1	0.1	-	43.7	0.1	-
	L23		867	63.0	5.1	5.1	0.6	7.2	8.1	0.6	-	7.8	-	-	-	2.5	-	-
	L1		53	6.0	0.3	-	16.0	5.0	5.8	12.9	8.3	1.6	0.1	0.1	-	43.7	0.1	-
B4		5.5	2359	8.0	0.7	-	15.1	5.2	5.8	14.8	7.3	1.5	0.1	-	-	41.6	-	-
NB4		1.5	2636	3.7	0.3	-	16.3	5.1	5.6	14.8	7.9	1.1	0.1	-	-	45.0	-	-
P5(L23)	L5	4.9	3971	49.9	2.0	-	3.6	2.2	8.2	1.0	-	12.7	1.1	2.1	12.5	2.5	2.2	-
	L4		198	4.0	0.1	-	17.4	5.4	6.0	9.5	8.4	1.1	0.1	-	-	48.0	-	-
	L23		413	62.9	5.1	5.1	0.6	7.2	8.1	0.6	-	7.8	-	-	-	2.4	-	-
	L1		12	97.7	1.7	-	-	0.3	1.5	0.1	-	1.1	-	-	-	-	-	-
P5(L56)	L5	1.3	4588	49.3	1.8	-	3.6	2.2	8.1	0.9	-	12.5	1.3	2.0	13.0	2.5	2.8	-
	L4		666	4.0	0.1	-	17.4	5.4	6.0	9.5	8.4	1.1	0.1	-	-	48.0	-	-
	L23		1368	63.0	5.1	5.1	0.6	7.2	8.1	0.6	-	7.8	-	-	-	2.5	-	-
	L1		375	95.6	1.7	-	-	0.3	1.5	0.1	-	1.1	-	-	-	-	-	-
B5		0.6	2744	49.5	2.5	-	3.6	2.2	8.1	1.2	-	12.6	1.1	2.3	12.3	2.5	2.2	-
NB5		0.8	2744	49.5	2.5	-	3.6	2.2	8.1	1.2	-	12.6	1.1	2.3	12.3	2.5	2.2	-
P6(L4)	L6	13.8	1326	6.1	0.3	-	1.8	2.1	3.1	0.2	-	0.2	12.0	0.9	-	2.9	32.4	38.1
	L5		979	51.0	0.9	-	3.7	2.3	8.4	0.6	-	13.0	1.1	1.6	12.7	2.5	2.3	-
	L4		1344	4.0	0.1	-	17.4	5.4	6.0	9.5	8.4	1.1	0.1	-	-	48.0	-	-
	L23		121	62.9	5.1	5.1	0.6	7.2	8.1	0.6	-	7.8	-	-	-	2.4	-	-
P6(L56)	L6	4.6	2264	4.0	0.1	-	17.4	5.4	6.0	9.5	8.4	1.1	0.1	-	-	48.0	-	-
	L5		236	51.0	0.9	-	3.7	2.3	8.4	0.6	-	13.0	1.1	1.6	12.7	2.5	2.3	-
	L4		171	4.0	0.1	-	17.4	5.4	6.0	9.4	8.4	1.1	0.1	-	-	47.8	-	-
	L23		286	63.1	5.1	5.1	0.6	7.2	8.1	0.6	-	7.8	-	-	-	2.5	-	-
	L1		4	97.7	1.7	-	-	0.3	1.5	0.1	-	1.1	-	-	-	-	-	-
B6		2.0	1310	6.1	0.3	-	1.8	2.1	3.1	0.2	-	0.2	12.0	0.9	-	2.9	32.4	38.1

Table A.1: Model composition. Neuron population sizes are given as percentage of total model size. The maximum number of synapses received by a postsynaptic neuron is specified per-layer for pyramidal neurons, whose apical dendrites span several layers. The proportions of these synapses made by each presynaptic neuron group are given in percentages of these maximal synapse numbers. Neurons in the slice model receive fewer than the maximum number of possible synapses because of the effects of slice cutting. Adapted from (Binzegger et al., 2004), with long-range connections removed.

Postsynaptic neurons \ Presynaptic neurons	P23	B23	NB23	SS4(L4)	SS4(L23)	P4	B4	NB4	P5(L23)	P5(L5/6)	B5	NB5	P6(L4)	P6(L5/6)	B6
P23	0.020	0.126	0.001	0.356	0.036	0.073	1.080	-	0.004	-	-	-	0.047	-	-
B23	0.560	0.026	0.001	0.701	0.078	0.161	0.228	-	0.074	-	-	-	0.159	-	-
NB23	0.408	0.069	0.014	0.872	0.085	0.173	0.581	-	0.159	-	-	-	0.178	-	-
SS4(L4)	0.001	0.043	-	0.067	0.092	0.061	0.010	0.003	0.069	0.069	-	-	0.004	-	-
SS4(L23)	0.001	0.043	-	0.067	0.092	0.061	0.011	0.003	0.069	0.069	-	-	0.004	-	-
P4	0.001	0.043	0.008	0.067	0.092	0.061	0.014	0.001	0.069	0.069	-	-	0.004	0.004	-
B4	0.101	0.098	-	0.627	0.627	0.627	0.025	0.003	0.841	0.841	-	-	0.062	0.062	-
NB4	0.139	0.244	-	0.318	0.318	0.318	0.068	0.013	1.058	1.058	-	-	0.052	-	-
P5(L23)	0.037	0.188	0.004	0.091	0.082	0.050	0.341	0.005	0.079	0.471	0.459	0.003	0.032	0.093	-
P5(L5/6)	0.037	0.188	0.004	0.091	0.082	0.050	0.289	0.005	0.062	0.335	0.416	0.003	0.032	0.093	-
B5	0.083	0.098	-	0.274	0.273	0.151	0.191	-	0.910	3.966	0.166	0.003	0.342	0.207	-
NB5	0.064	0.244	-	0.331	0.422	0.196	0.521	-	0.603	2.596	0.166	0.014	0.359	0.657	-
P6(L4)	0.003	1.045	0.015	0.137	0.145	0.095	0.226	0.001	0.084	0.055	0.293	0.004	0.064	0.062	0.075
P6(L5/6)	0.003	1.045	0.015	0.137	0.145	0.095	0.978	0.016	0.201	0.055	0.293	0.004	0.064	0.062	0.048
B6	0.123	0.140	-	0.274	0.273	0.151	0.193	-	0.091	0.091	0.021	-	1.105	0.768	0.015

Table A.2: Synaptic weights, in nS.

Postsynaptic neurons \ Presynaptic neurons	P23, P4	SS	P5	P6	B	NB
P23, P4	3,6,7,8	6-8	4,5	4	1,2,6	3-5,7,8
SS4	3,4,6,7	3,4,6,7	3,4,6,7	3,4,6,7	1,2,5	3,4,6,7
P5	2-9	2-5	2-5,7-9	2-5,7-9	1,2,7	3-5,7,8
P6	2-9	2,4,5	2-9	2-9	1,2,7	3-5,7,8
B	3,4,6,7	3,4,6,7	3,4,6,7	3,4,6,7	3,4,6,7	3,4,6,7
NB	3,4,6,7	3,4,6,7	3,4,6,7	3,4,6,7	3,4,6,7	3,4,6,7

Table A.3: Compartment IDs in each postsynaptic group that presynaptic neurons connect onto. Based on (Traub et al., 2005).

Appendix B

Experimental procedures *in vitro*

Here we summarise the experimental approach used to collect the comparison slice data presented in Chapters 4 and 5. Methods provided by Matt Ainsworth and Mark Cunningham. Experiments performed by Matt Ainsworth.

All experiments were carried out in accordance with the European Communities Council Directive 1986 (86/609/EEC), the US National Institutes of Health Guidelines for the Care and Use of Animals for Experimental Procedures, and the UK Animals Scientific Procedures Act.

B.1 Surgical preparation

The monkey (Macaca mulatta, male, 8 years old) used in this study was subject to experiments *in vivo* involving extracellular recording of neural activity and local drug application (iontophoresis). All tissue samples used in this study were taken from intact brain areas that were not the subject of studies performed before tissue extraction. Tissue extraction was performed under general anaesthesia, which was maintained over the course of four days. For the anaesthesia the animal was initially sedated with a 0.1 ml/kg ketamine intra-muscular injection (100mg/ml). Thereafter, bolus injections of propofol were administered intravenously to allow for tracheotomy and placement of catheters for measuring intra-arterial and central venous blood pressure. During surgery, anaesthesia was maintained by gaseous anaesthetic (2.5-3.9% sevoflurane) combined with continuous intravenous application of an opioid analgesic (Alfentanil, 120μg/kg/h), a glucocorticoid (Methylprednisolone, 5.4mg/kg/h) and saline (50ml/h). The animal's rectal temperature, heart

rate, blood oxygenation and expired CO2 were monitored continuously during anaesthesia.

B.2 Slice preparation

Neocortical slices of 450 micron thickness were prepared from tissue resected from macaque temporal cortex. Following resection, cortical samples were immediately placed in ice-cold sucrose artificial cerebrospinal fluid (ACSF). Neocortical slices containing all layers were cut at 450 micron (Microm HM 650V), incubated at room temperature for 20-30 minutes, then transferred to a standard interface recording chamber at 34-36°C perfused with oxygenated ACSF. Persistent gamma frequency oscillations were induced by the application of kainate (400nM - 800 nM) to the circulating ACSF and were deemed stable if there was no change to frequency or power after 1 hour. LFP recordings were taken using multichannel 10x10 silicon electrodes with an inter-electrode distance of 400 microns (Utah array, Blackrock Microsystems, Salt Lake City, UT, USA). Time series were digitally sampled at 10 kHz.

Appendix C

VERTEX user guide

A key outcome of this thesis was the development of the Virtual Electrode Recording Tool for EXtracellular potentials (VERTEX), described in Chapter 3. In this appendix, we provide a user manual and example code for simulating the layer 2/3 model studied in Chapter 5. VERTEX is still under active development, so the underlying code is likely to change in the future. While we do not intend to make significant changes to the top-level interface, small changes may render sections of this guide inaccurate in future. Any changes to this guide will be made available at Marcus Kaiser's web-site www.biological-networks.org along with the latest version of VERTEX.

C.1 Introduction

VERTEX is a Matlab tool designed to facilitate the simulation of extracellular potentials generated by activity in spiking neural networks; in particular, spatially-organised networks containing thousands or hundreds of thousands of neurons. VERTEX's interface and model specification options were designed with this particular task in mind. It is therefore less flexible than other neural simulators (e.g. NEURON, NEST, Brian, GENESIS, Moose), but the limited scope has allowed us to simplify the user interface so that a simulation can be specified simply by setting some parameters and run using a few function calls. The parameters are divided into five categories: neuron group properties, ConnectionParams, tissue properties, recording settings, and simulation settings. The parameters associated with each category are specified in Matlab structures or structure arrays.

C.2 Installation

To install VERTEX, first download the code from the link provided at www.biological-networks.org and unzip the archive to the directory you want VERTEX to reside in. To start using VERTEX in Matlab, you simply need to add this directory and its sub-directories to Matlab's path (right click on the VERTEX directory within Matlab and choose 'Add to Path' -> 'Selected Folders and Subfolders'). Currently only one function in VERTEX is implemented in C using Matlab's Mex interface (though this number is likely to increase in future, we will always provide pure Matlab functions for people that have trouble compiling the C code). If you can't compile this file, VERTEX will still run, but you won't be able to allow multiple synapses between two neurons (see below). To compile the mex file, go into your VERTEX directory, then into the 'vertex_mex' directory. If you have Matlab set up properly with your C compiler, you should simply be able to type `mex multiSynapse.cpp` and Matlab will compile the function, making it available to VERTEX. For instructions on how to set up a C compiler with Matlab, see the MathWorks web-site.

You should now be able to use VERTEX to create and run simulations. In the next section, we walk through setup of a simple, single neuron population network containing randomly firing reduced layer 2/3 pyramidal neurons.

C.3 Simple simulation walkthrough

C.3.1 Tissue parameters

In this example, we walk through the different parameter settings for creating a single group of neurons that synapses with itself. We start with the structure array holding the overall information about the piece of brain tissue we are modelling, which we call `TissueParams`. The order in which parameters are specified does not matter to VERTEX, so we proceed in what seems to us like a reasonable order, starting with the size of the model.

```
TissueParams.X = 2500;
TissueParams.Y = 400;
TissueParams.Z = 200;
TissueParams.neuronDensity = 50000;
```

`X`, `Y` and `Z` specify the dimensions of the model, in micrometres. The model space in this case is cubic, as in the brain slice models described in this thesis. Alternatively, if we do not specify `X` or `Y`, but instead specify `R`, then VERTEX will create a cylindrical model with radius `R` micrometres and depth `Z` micrometres. `neuronDensity` gives the overall density of neurons in the model, in neurons per cubic millimetre. The number of neurons in the model is then calculated automatically (in this case there will be $2.5 \times 0.4 \times 0.2 \times 50,000 = 10,000$ neurons in the model).

Next we specify some more parameters concerning the spatial layout of the model:

```
TissueParams.numLayers = 1;
TissueParams.layerBoundaryArr = [200,  0];
TissueParams.numStrips = 10;
```

`numLayers` is the total number of layers we want to create in the model. In this simple case, we just have the one layer. As we assume the layer boundaries are defined as x-y planes, we can set the boundaries by specifying their z-depths. `layerBoundaryArr` contains a list of these depths, going from the top of the highest layer to the bottom of the lowest one. The length of `layerBoundaryArr` should be equal to `numLayers+1`. Finally, we introduce the concept of strips. When setting up the neuron positions in a cuboid model, VERTEX will by default position each neuron's soma randomly within the boundaries of its containing layer. This often does not matter, but sometimes it is useful to know more precisely where the neurons are positioned - when plotting spike trains to reveal spatial variation in firing patterns, for example. A cuboid model can be divided into two or more spatial 'strips', which are ordered from left to right on the x-axis. Neurons are then positioned by strip so that the lowest neuron IDs are at the strip to the farthest left, and the highest neuron IDs are at the strip to the farthest right (see Figure C.1). The width of each strip is the width of the model, `X`, divided by the number of strips. Changing `numStrips` does not affect the model's dynamics as such; it affects the precision of the placement of neurons along the x-axis.

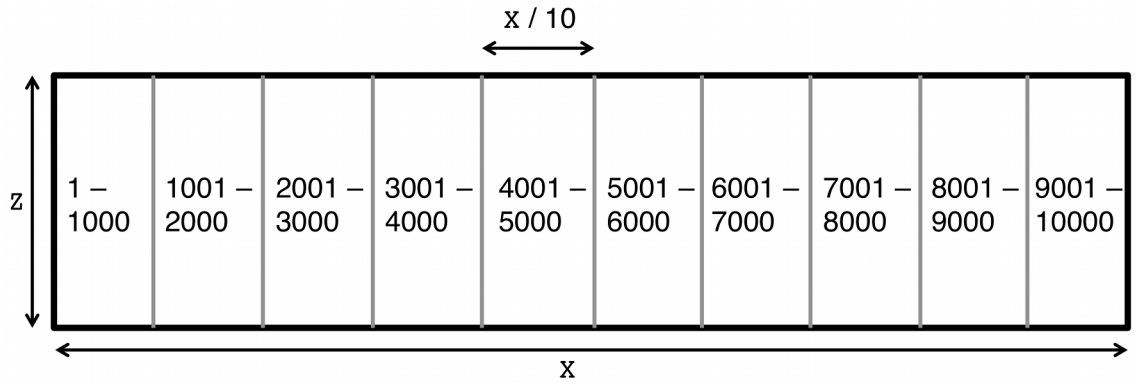

Figure C.1: Arrangement of strips for a cubic model, showing a model with 10,000 neurons and 10 strips. Ids of the neurons contained within the strips are shown in each strip location.

`TissueParams` contains two final parameters:

```
TissueParams.tissueConductivity = 0.3;
TissueParams.maxZOverlap = [-1, -1];
```

`tissueConductivity` sets the conductivity of the extracellular medium, in Siemens per metre. This parameter is only required when using simulated extracellular electrodes. `maxZOverlap` sets the maximum distance outside the model's upper and lower z-boundaries that a neuron's dendrites can extend. As neurons are positioned according to their soma compartment, if the soma is randomly placed close to the model's boundaries, then its dendritic compartments will extend outside these boundaries. In the present model, we set the model z-depth to be 200 microns, so pyramidal neuron apical dendrites are bound to extend above the model space. In the current case this is desirable, as we set the 200 micron depth to be the depth of the soma-layer rather than meaning it to be the depth of the whole layer. Therefore we set `maxZOverlap` to be -1, meaning that we do not want to set a maximum z-overlap length. In the 15 population cortical slice model described in chapter 4, though, the layer boundaries are meant to represent the entire depth of the layer and the full size of the model. In that case, we set `maxZOverlap` to be `[0, 100]`. This prevented the uppermost pyramidal neuron dendrite compartments from above the top of layer 1, but allowed dendritic compartments at the bottom of layer 6 to extend below the model by up to 100 microns into the "white matter".

C.3.2 Neuron group parameters

We have now finished specifying the `TissueParams` structure, so now we specify the neuron group information in a structure array we call `NeuronParams`. We store the parameters for our only neuron group in `NeuronParams(1)`, in case we want to add further groups at a later date. This neuron group is henceforth referred to as neuron group 1.

```
NeuronParams(1).modelProportion = 1;
NeuronParams(1).somaLayer = 1;
NeuronParams(1).neuronModel = 'passive';
```

The first line defines the proportion of the total model size made up by neuron group 1. In this case we are only modelling a single neuron group, so the whole model is made up of neurons in this group. The second line specifies which layer of the model this neuron group is in. The third line specifies neuron group 1 as having passive cell dynamics. These kinds of groups can be useful for simulating the LFP, but will not generate any spikes in the simulation. Instead, we will specify group 1 to consist of neurons randomly firing with Poisson statistics, and a firing rate of 5 Hz:

```
NeuronParams(1).neuronModel = 'poisson';
NeuronParams(1).firingRate = 5;
```

Next we specify the number of compartments per neuron we want.

```
NeuronParams(1).numCompartments = 1;
```

If we specify 1 compartment, then VERTEX will treat this group as consisting of point neurons. No further parameters would be required in the `NeuronParams` structure array in this case. This is useful for specifying groups of "external" populations that provide external input to the simulation, but do not contribute to the LFP, for example. However, as we want to simulate the LFP in this group, we need

to use compartmental neurons. We will use the layer 2/3 pyramidal neuron models described in Chapter 2, which contain 8 compartments.

```
NeuronParams(1).numCompartments = 8;
NeuronParams(1).compartmentParentArr = [0, 1, 2, 2, 4, 1, 6, 6];
NeuronParams(1).compartmentDiameterArr = [29.8, 3.75, 1.91, 2.81, ...
                                          2.69, 2.62, 1.69, 1.69];
NeuronParams(1).compartmentXPositionMat = [0, 0; 0, 0; 0, 124; 0, 0; 0, 0; ...
                                           0, 0; 0, -139; 0, 139];
NeuronParams(1).compartmentYPositionMat = [0, 0; 0, 0; 0, 0; 0, 0; 0, 0; 0, 0; ...
                                           0, 0; 0, 0; 0, 0];
NeuronParams(1).compartmentZPositionMat = [-13, 0; 0, 48; 48, 48; 48, 193; ...
                                           193, 330; -13, -53; -53, -139; ...
                                           -53, -139];
```

`compartmentParentArr` gives the compartmental structure. As we have assumed our neurons will always be binary trees (see Chapter 3), each compartment will have a maximum of one parent compartment - and no parent for the soma. `compartmentParentArr` lists the parent compartment ID at the index of each child compartment; compartment ID 1 is the soma so has no parent compartment and is set to 0. `compartmentDiameterArr` gives the diameter of each compartment in microns, again with the indices of the array representing the compartment IDs. The `compartment[XYZ]PositionMat` values specify the start and end coordinates of each compartment, presuming the soma compartment is positioned at (0,0,0). During network construction, these values are used to calculate the position of each compartment of each neuron after random positioning and rotation. They are also used to calculate the length of each compartment (though a `compartmentLengthArr` list containing the length of each compartment as defined by the user can be specified to override the automatic calculations). Row indices represent compartment IDs, column 1 holds compartment start coordinates and column 2 holds column end coordinates.

Next we specify the neurons' passive properties:

```
NeuronParams(1).C = 1.0*2.96;
NeuronParams(1).R_M = 20000/2.96;
NeuronParams(1).R_A = 150;
NeuronParams(1).E_leak = -70;
```

`C` is the specific membrane capacitance in μFcm^{-2}, `R_M` is the specific membrane

resistance in Ωcm^{-2}, `R_A` is the axial resistance in Ωcm and `E_leak` is the membrane leak conductance reversal potential in mV. If we were using the AdEx model (by specifying `NeuronParams(1).neuronModel = 'adex';`) then we would also need to specify the other AdEx parameters (`delta_t`, `V_t`, `v_cutoff`, `tau_w`, `a`, `b`, `v_reset`).

The next code listing, in italics, is entirely optional, but useful when specifying the model ConnectionParams. We specify labels for the different compartments in the model, so that we don't have to remember explicitly the numbers of the compartments that comprise different sections of the neuron. The name of each parameter is chosen to be descriptive, but could be chosen arbitrarily.

```
NeuronParams(1).somaID = 1;
NeuronParams(1).basalID = [6, 7, 8];
NeuronParams(1).proximalID = [2, 6];
NeuronParams(1).distalID = [7, 8];
NeuronParams(1).obliqueID = 3;
NeuronParams(1).apicalID = [2 3 4 5];
NeuronParams(1).trunkID = 2;
NeuronParams(1).tuftID = 5;
```

As we are only specifying one neuron group, we are now finished with the `NeuronParams` structure array (further groups would have their parameters specified in `NeuronParams(2)`, `NeuronParams(3)` etc.).

C.3.3 Input parameters

In the current example, neurons fire randomly so synaptic input does not affect their spiking output at all. However, in most simulations we will be using neurons with dynamics that cause them to spike, and behaviour will emerge from the dynamics of this spiking and the network connectivity. We therefore need some way to drive the neurons to start spiking in the first place - input currents or conductances. We can specify these in another structure called `Input`. Each neuron group in the `NeuronParams` structure array can have an Input structure array attached to it to define inputs to the population.

```
NeuronParams(1).Input(1).inputType = 'I_ou';
NeuronParams(1).Input(1).meanInput = 330;
NeuronParams(1).Input(1).tau = 2;
NeuronParams(1).Input(1).std = 90;
NeuronParams(1).Input(1).startTime = 0;
NeuronParams(1).Input(1).compartments = [7, 8];
```

In this example, we have one Input of type `'I_ou'`, which adds an independent random current (`I`) defined as an Ornstein-Uhlenbeck process (`ou`) to each neuron in the group. We set the mean input current to 330 pA, the time constant `tau` to 2 ms and the standard deviation `std` to 90 pA. We also tell VERTEX that we want this current to be applied starting at time `startTime = 0`, so right from the start of the simulation (we can have currents starting after set amounts of time by changing this parameter). Finally, we specify which compartments of the neuron the currents should be delivered to - in this example, just the two basal dendrite compartments. When specifying multiple compartments like this, the same current will be applied to each compartment, scaled down by the compartment's membrane area. To add independent currents to different compartments, you can specify additional inputs in the `Input` structure array. The Ornstein-Uhlenbeck process is the only currently implemented random process for stimulation, but a DC stimulus can be applied by setting the `std` parameter to 0. We also provide a conductance-based version: `inputType = 'g_ou'`. The parameters to set for the conductance-based input are identical as those for the current-based input, except that `meanInput` and `std` are now in nanoSiemens, and you will additionally need to specify a reversal potential in a parameter called `E_rev`.

In our simple network example, adding inputs to the neuron group will not change the network dynamics, but will alter the LFP.

C.3.4 Connectivity parameters

VERTEX treats spikes as discrete events that trigger some postsynaptic event at targeted neurons. Connectivity is specified probabilistically in terms of connections between groups of neurons. We specify the relevant parameters in a structure array called `ConnectionParams`.

```
ConnectionParams(1).prePost = 'pre';
```

First we need to establish whether we want to connect the neurons using a presynaptic or postsynaptic perspective (see Chapter 3, section 3.3.2). The connection algorithms establishes connections one neuron at a time. If we choose a postsynaptic perspective, then each neuron has a set of presynaptic neurons selected that connect to it. If we choose a presynaptic perspective, then each neuron has a set of postsynaptic targets selected to connect to. The sets of presynaptic/postsynaptic neurons are chosen randomly, according to some distance rule (specified below). The effect of choosing one method over the other is that either the in-degree (postsynaptic perspective) or out-degree (presynaptic perspective) of each neuron can be controlled. Additionally, if a presynaptic perspective is chosen and the model is a cube, then the number of connections made by each presynaptic neuron will automatically be scaled down proportionally to the amount of its axonal arbor falling outside the model boundaries - see Chapter 3 for details.

In terms of parameter specification, choosing to set `prePost` to either `'pre'` or `'post'` has consequences for the interpretation of the remainder of the parameters specified in `ConnectionParams`. Most parameters are specified for connections involving two neuron groups, rather than just for one group. Parameters are therefore specified in cell arrays. When taking the postsynaptic perspective, we assume that the index of the structure array refers to the postsynaptic neuron group, and the index of the cell array refers to the presynaptic group. For example:

```
ConnectionParams(1).numConnectionsToOneFromAll{13} = 1000;
```

means that each neuron in postsynaptic group 1 receives 1000 connections from neurons in neurons in group 13. In our example model, though, we have only one population and specify a presynaptic perspective:

```
ConnectionParams(1).numConnectionsToAllFromOne{1} = 1000;
ConnectionParams(1).synapseTypes{1} = 'i_exp';
ConnectionParams(1).targetCompartments{1} = [NeuronParams(1).basalID, ...
                                             NeuronParams(1).apicalID];
ConnectionParams(1).weights{1} = 10;
ConnectionParams(1).tau{1} = 2;
```

In the above code, we first say that each presynaptic neuron in group 1 (index of the structure array) makes 1000 connections to postsynaptic neurons in group 1 (index of the cell array). We specify that the type of these synapses is `'i_exp'` - single exponential, current-based synapses (the other currently available synapse type in VERTEX is `'g_exp'`, which are conductance-based single exponential synapses; in this case, a reversal potential for the synapse must be specified in the relevant `E_rev` parameter cell). `targetCompartments` lists which compartments on the postsynaptic neurons are allowed to be contacted by the presynaptic neurons (this is where the labels we gave to the compartment numbers comes in handy). Finally we specify the weight of these synapses - for current-based synapses in picoAmps, and for conductance-based in nanoSiemens - and the exponential decay time constant in milliseconds.

When a model contains more than one layer, we can specify the number of synapses on a per-layer as well as per-group basis. In the 5 layer model described in chapter 4, for example, each element of the `numConnectionsToAllFromOne` cell array contains an array of length 5, specifying the number of connections to be made in each layer from a presynaptic neuron to the postsynaptic group. For example:

```
ConnectionParams(1).numConnectionsToAllFromOne{13} = [100, 200, 500, 200, 50];
```

In this example, each neuron in group 1 makes 100 connections with neurons in group 13 in layer 1, 200 connections in layer 2, 500 in layer 3, 200 in layer 4 and 50 in layer 5. This is useful in particular for pyramidal neurons spanning several layers.

Finally, we need to specify the axonal arbour properties.

```
ConnectionParams(1).axonArborSpatialModel = 'gaussian';
ConnectionParams(1).axonArborRadius = 250;
ConnectionParams(1).axonArborLimit = 500;
```

The axon arbour of each neuron is defined as a 2D probability distribution in the x-y plane centred at the presynaptic neuron, and defines how the probability of making a connection to a postsynaptic neuron varies in space. Specifying the `axonArborSpatialModel` to be `'gaussian'` means that the connection probability decays away from a presynaptic neuron as a 2D gaussian with standard deviation given by the `axonArborRadius` parameter. If we instead set `axonArborSpatialModel` to be `'constant'`, the `axonArborRadius` parameter would represent the maximum extent of the arbor and the connection probability would be constant within this radius. These are the two currently implemented connectivity profiles; we intend to add more complex profiles (e.g. non-isotropic gaussians, patchy projections) in the future. `axonArborLimit` is an optional parameter for gaussian arbors, and specifies a cutoff point beyond which no connections can be made. If this parameter is not set, then neurons have no limit on their maximum connection distance.

Axon arbor radii and limits can be specified per-layer in the same way as the number of connections; for example in a 5 layer model we could specify:

```
ConnectionParams(1).axonArborRadius = [100 250 200 500 100];
ConnectionParams(1).axonArborLimit = [200 500 400 2000 200];
```

C.3.5 Recording parameters

In order to analyse the output of a model, we need to be able to record the variables over the time course of the simulation. We specify what we want to record in the recording parameters structure. VERTEX will save the spike trains of all neurons by default, but can also record individual neuronal membrane potentials and LFPs.

```
RecordingSettings.saveDir = '~/VERTEX_results/';
RecordingSettings.LFP = true;
[meaX, meaY, meaZ] = meshgrid(0:100:2500, 200, 650:100:-50);
RecordingSettings.meaXpositions = meaX;
RecordingSettings.meaYpositions = meaY;
RecordingSettings.meaZpositions = meaZ;
RecordingSettings.minDistToElectrodeTip = 20;
```

First we set the directory in which we want to save the files that VERTEX

generates. Then we set `LFP` to `true`, so VERTEX knows we want to record LFPs. We then use Matlab's `meshgrid()` function to calculate x, y and z coordinates (in microns) of our multi-electrode array. VERTEX expects electrode coordinates in the format output by `meshgrid`, but it isn't necessary to use the `meshgrid` function to specify the electrode positions - they can be placed arbitrarily. Model coordinates in VERTEX use standard x, y and z axis directions, so we specify the z-coordinates of the MEA in descending order so that they are numbered from top to bottom as in the experimental MEA in Chapter 4. This example produces a 2.5 mm by 0.7 mm MEA with 208 electrodes, at a constant y-coordinate of 0.2 mm, with an inter-electrode spacing of 0.1 mm. Finally we set `minDistToElectrodeTip` - this value ensures that no neuronal compartment can be positioned too close to an electrode tip and therefore unrealistically dominate the LFP. We choose 20 microns, which is the default value if this parameter is not set by the user.

We can also choose which neurons we want to record the membrane potential from, if any. We provide this in a list of neuron IDs. While the IDs are only generated during model initialisation, we know roughly where in the model space the neuron of each ID will be because of the method of using strips to constrain the x-coordinates of the neurons. We can sample membrane potentials of one neuron from each strip, for example, to look at differences in membrane potentials across the space of the model. Referring to Figure C.1, we will specify:

```
RecordingSettings.v_m = 500:1000:9500;
```

to tell VERTEX to record from neuron IDs 500, 1500, 2500, ... 9500, ensuring that we record from one neuron in each strip. The membrane potential is always recorded at the soma compartment.

Next we set the sample rate and the maximum simulation time between saves:

```
RecordingSettings.maxRecTime = 200;
RecordingSettings.sampleRate = 1000;
```

`maxRecTime` specifies, in milliseconds, the amount of time that VERTEX records

for before saving the recordings to disk and starting recording again. This chunks the recordings into files of length `maxRecTime` × `sampleRate` / 1000, and ensures that for large MEAs the file sizes and memory usage do not grow too large. We have found that setting this to 200 ms provides a good compromise between memory usage and slowing the simulation down due to excessive file save operations, though the optimal value will depend on your model details. In smaller models with few electrodes, it may be suitable to set this to be equal to the total simulation length. sampleRate is simply the rate at which VERTEX samples the LFP, in Hz. High values can severely slow simulation time in large models, or for large numbers of electrodes. 1000 Hz provides a good resolution for most LFP investigations (note: in the next section we will describe how to set the simulation timestep. If the simulation timestep and specified sample rate are incompatible, VERTEX will automatically sample at the closest (higher) frequency that is compatible with the timestep).

C.3.6 Simulation parameters

Finally we need to specify some overall simulation settings.

```
SimulationSettings.simulationTime = 1000;
SimulationSettings.timeStep = 0.03125;
SimulationSettings.parallelSim = true;
SimulationSettings.poolSize = 4;
SimulationSettings.profileName = 'local';
```

We set the total simulation time to 1000 ms, and the global integration timestep to 0.03125 ms. We then tell VERTEX to create a parallel simulation by setting `parallelSim` to true, setting the Matlab pool size, and defining the Matlab parallel profile to use. If `poolSize` and `profileName` are not set, then VERTEX will ask Matlab to start a parallel job with the default settings.

C.3.7 Generating the network

Now that we have set all the required parameters for our simulation in the relevant structures, we can generate the network.

```
Model_1 = initialiseNetwork(TissueParams, NeuronParams, ...
            ConnectionParams, RecordingSettings, SimulationSettings);
```

The `initialiseNetwork()` function performs a number of operations. Firstly it sets up the parallel environment if a Matlab parallel pool is not already running. Secondly, it creates the neuron groups, positions them in space, and if necessary distributes them between parallel processes. Thirdly, it calculates the model connectivity based on the connection statistics and neuron positions. Finally, if recording the LFP, it pre-calculates the constants required for the line-source extracellular potential calculation based on the neuron and electrode positions. `initialiseNetwork()` returns all this in an Matlab object, with class `VertexNetwork` (you don't need to know about object-oriented programming in Matlab to use VERTEX). In our example we call this `Model_1`. After the model is initialised, we can save it to disk using the `vertex_save` function:

```
vertex_save('~/vertex_models/Model_1.mat', Model_1);
```

and similarly load it back into Matlab using `vertex_load(path_to_model)`. These functions deal with both serial and parallel simulation setups (though so far parallel saving has only been tested in the case of a parallel job running on a single multicore machine). In the parallel case, the standard Matlab Parallel Computing Toolbox `psave()` function is used, which saves each lab's data in a separate file, numbered with the ID of the lab it resided on. `vertex_load()` will deal with this case or the serially saved case automatically.

C.3.8 Running a simulation

Now that we have generated the network, we can run the model to simulate the dynamics in the network.

```
runSimulation(Model_1);
```

This function will run the dynamics in the simulation according to the specified parameters and previously generated network, storing the variables that have been specified to record from in the location set in `RecordingSettings`. The timestep will be printed every few milliseconds so you can keep track of how long the simulation has left to run. Once the simulation has run (this may take some time depending on your computer and the size of the model), you can load the results for analysis.

C.3.9 Loading the results

VERTEX provides a script called `loadSimulationResults.m` which can be run to load and plot the results of a simulation. The user needs to provide the script with the directory in which the simulation results were saved, and the rest of the loading is then handled automatically. Running the script in its entirety will produce simple plots of each of the recorded variables.

Spike times are loaded into an N by 2 matrix called `spikes`, where N is the number of spikes. The first column of this matrix contains the ID of the neuron that fired, and the second column contains the time at which the neuron fired. The user can write their own scripts to plot this data, or can use VERTEX's `plotSpikeRaster(spikes)` function.

LFPs are loaded into a N_E by r_s matrix called `LFP`, where N_E is the number of electrodes and r_s is the sample rate used. To plot the LFP from electrode 1, for example, you would use `plot(LFP(1, :))`.

Membrane potentials are loaded into a N_v by r_s matrix called `v_m`, where N_v is the number of neurons recorded from. The ordering of the rows in this matrix is the same as the order of the neuron IDs specified in the `RecordingSettings` structure.

C.4 List of physical units used in VERTEX

VERTEX uses a consistent set of units, which we list below in Table C.1. Any time a physical quantity is specified as a parameter, it will be in the units specified in this table.

Physical quantity	**Unit used**	**Example**
Electric potential	mV	membrane potential v
Conductance	nS	synaptic conductance g_{syn}
Current	pA	synaptic current I_{syn}
Specific resistance	$\Omega \cdot cm^2$	Specific membrane resistance R_M
Specific capacitance	$\mu\mathrm{F} \cdot cm^{-2}$	Specific membrane capacitance C_M
Longitudinal resistance	$\Omega \cdot \mathrm{cm}$	Intracellular axial resistance R_A
Conductivity	$\mathrm{S} \cdot \mathrm{m}^{-1}$	Extracellular conductivity σ
Time	ms	Synaptic time constant τ_{syn}
Length	μm	Model width X

Table C.1: Units used in VERTEX