# Single unit action potential recordings in humans: implications for epilepsy

**Edward M. Merricks**

Thesis submitted for the degree of

Doctor of Philosophy

Institute of Neuroscience

Newcastle University

October 2015

## Abstract

Spike-sorting algorithms have been used to identify the firing patterns of isolated neurons ('single units') from implanted electrode recordings in patients undergoing assessment for epilepsy surgery, but we do not know their potential for providing helpful clinical information. It is important to characterize both the stability of these recordings and also their context. I therefore analysed microelectrode array recordings from four patients undergoing video-telemetry monitoring for surgical evaluation of focal neocortical epilepsies. I investigated whether units could be followed reliably over prolonged periods during the initial days post-implantation in humans. Unit specific features showed stability over at least 48 hours, including across multiple seizures.

A critical consideration regarding the clinical information in these recordings is where the units are located with respect to the focus of pathology. Recent analyses have demonstrated the importance of considering seizure activity in terms of two distinct territories: the ictal core and penumbra. The pathological information in these two areas is likely very different. I isolated unit recordings from several hundred neurons across the four patients in the peri-ictal period, with reference to these regions. A key finding was that in the penumbra, spike stereotypy was maintained even during the seizure, with little alteration to population firing rate. In contrast, within ictal core territories, regions characterized by intense hypersynchronous firing, spike sorting failed as the units were incorporated into the seizure activity.

Recovery of spike shape was rapid following seizure termination, and the mean firing rate returned to pre-ictal levels in the first few minutes. However, the spatiotemporal activity of units in the ictal core displayed significant alterations to firing patterns, lasting at least 30 minutes after seizure termination. These observations lay the foundation for future investigations of how these recordings may inform clinical practice.

## Acknowledgements

First and foremost, I would like to extend my gratitude to Dr. Andrew J. Trevelyan, who has been a fantastic supervisor throughout. Andy introduced me to the world of scientific research, and has unendingly provided the invaluable guidance, support, and encouragement needed through the whole project. It is thanks to Andy that I found my passion within neuroscience.

Thanks also to the members of "Trev Lab": Dr. Rolando Berlinguer-Palmini; Dr. Ryley Parrish; Dr. Hannah Alfonsa; Neela Codadu; Partow Yazdani; Christoforos Papasavvas; and Emma Craddock. You have all played a great part in my continued enjoyment throughout the process of creating this, over the past three years. Indeed, a special thank you to Neela, for our in-depth discussions; even if we never got around to actually performing the copious intended experiments that resulted from them. Not to mention your wondrous talents in front of a spice rack.

I would especially like to thank Dr. Catherine A. Schevon, who has been almost a foster second supervisor across the pond over the last three years, and through whom all of this has been possible. Thank you, indeed, to the various members of the clinical and research team at Columbia University Medical Center for greatly valued discussion and feedback throughout. Thanks to Dr. Elliot H. Smith, who has tirelessly taken the time to run Matlab scripts for me, which I have sent over at short notice when I have realised that I need yet more data. And particularly, thank you to the study patients, for their selfless participation in these projects.

I would like to acknowledge the Wellcome Trust, first for funding this project, and, more generally, for their willingness to fund students of far-fetched undergraduate backgrounds wishing to make the jump into neuroscience.

Thanks, of course, to my family, without whom I would likely not have found this path, and for their continued encouragement, and thank you to Elishia for putting up with me through the whole process.

# Table of Contents

# List of Figures

# List of Tables

# Chapter 1. Introduction

The human brain is comprised of upwards of 100 billion neurons (Buzsàki, 2006). It is through the interactions of these neurons that sensations, perceptions, and resultant consciousness arise, with each individual cell making hundreds to thousands of connections to others (Shepherd, 2004). These interactions are maintained and controlled through differing downstream effects of separate neuronal types. A subset of neurons provide inhibition onto others, resulting in those neurons being less likely to fire an action potential at that moment in time, while other neurons excite their downstream synaptic connections, causing an increase in the likelihood of firing. In the cortex, these inhibitory and excitatory cells are at a distribution of roughly 20% and 80% respectively (Hendry et al., 1987).

An oft used description of the method underlying cortical processing is the maintenance of a delicate balance between the inhibition and excitation (Nelson & Turrigiano, 1998). In reality, this interaction between inhibition and excitation is considerably more complex than a linear balance, with the presynaptic cell providing different types of inhibition depending upon where the synapse is formed on the postsynaptic cell, and the ratio between excitation and inhibition is highly dynamic (Isaacson & Scanziani, 2011; Pouille et al., 2013). Somatostatin and parvalbumin containing inhibitory interneurons target the dendritic processes and somata respectively (Megías et al., 2001), and in doing so provide subtractive and divisive inhibition to the excitatory current in the cell (Wilson et al., 2012). As a result, very fine tuning of cortical activity may be achieved through the weights of influence that presynaptic cells have upon a neuron.

Increasingly complex processing can thus be achieved through the interplay of dendritic activity, where a single postsynaptic potential cannot alone cause or halt an action potential, but instead takes part in a weighted “conversation” as to the eventual activity of the cell. This delicate activity is combined with the properties of inhibition in the peri-somatic region, that has more direct ability to halt the onset of an action potential through shunting inhibition (Jack et al., 1975). Through these complex interactions over both time and space,

cortical processing is performed through the resultant timings of the neurons firing an action potential within both local and global networks (Varela et al., 2001). Therefore, through analyses of these firing patterns of individual neurons, we may probe the underlying mechanisms of cognition and consciousness.

## 1.1. Clinical epilepsy

Given the complex interactions between inhibition and excitation in the cortex, pathological states may arise as disruptions of the delicate interplay between the two. Indeed, abnormal occurrences of excessive or synchronous activity at the neuronal level can result in seizures. A diagnosis of epilepsy is typical when a patient has a long lasting predisposition to such seizures, and has had at least one seizure, as defined recently by the International League Against Epilepsy (ILAE; Fisher et al., 2014).

This predisposition to seizures is not rare; epilepsy is the fourth most common neurological disorder (Hauser et al., 1993), affecting more than 65 million people worldwide (Moshé et al., 2015). The risk for an individual developing epilepsy by age 50 is greater than 1%, increasing to higher than 3% by the age of 80 (Hesdorffer et al., 2011). This prevalence includes a complex range of seizure types, and potential aetiologies. Despite a long history of medical studies into epileptic seizures, including an in depth description as far back as the late 19$^{th}$ century by John Hughlings Jackson (Jackson, 1879), up to 70% of newly diagnosed cases have no known aetiology (Epilepsy Foundation of America, 1999). Moreover, up to 30% of epilepsy patients remain refractory to pharmacological intervention and continue to have seizures despite attempts at management with multiple drugs (Moshé et al., 2015).

### 1.1.1. *Ictal activity*

Seizures themselves are varied between patients, with discrepant outward manifestations depending on where the ictal activity is occurring. From the simplest standpoint, seizures can be separated into partial and generalized types (Seino, 2006), though partial seizures may secondarily generalize (Fisher et al., 2014). Partial seizures are, by definition, focal, incorporating a limited region of the brain, be it cortical or subcortical. In generalized seizures on the other hand, the ictal activity incorporates large areas of the brain, though in recent years, whether or not this is truly a global incorporation has been questioned, with

the evidence leaning towards, rather, a distributed network of focal ictal activity (Rektor et al., 2009; Schindler et al., 2007; Holmes et al., 2004).

In all instances, the enduring understanding of ictal activity, in the regions of cortex that have been recruited to the seizure, is of hypersynchronous bursts of firing across the majority of the population (Traub & Wong, 1982), and of paroxysmal depolarising shifts (PDSs) at the level of individual neurons (Matsumoto & Marsan, 1964; Steriade et al., 1998). This intense, hypersynchronous activity projects a large glutamatergic barrage upon the neurons that are post-synaptic to those involved in the ictal activity. This large synaptic barrage is hypothesized to be either successfully quelled through feed forward inhibition, or for the inhibition to give way, allowing downstream tissue to become incorporated also, thus leading to the propagation of seizures through the brain (Trevelyan et al., 2006; Trevelyan et al., 2007; Trevelyan & Schevon, 2013). These underlying mechanisms are discussed further throughout section 1.2.

### 1.1.2. *Comorbidities in epilepsy*

The seizure itself is, of course, the defining feature of epilepsy, but it is not where the clinical symptoms end. During the time between seizures (the "inter-ictal" period), bursts of epileptiform activity often occur, which can have a severe effect on cognition (Kleen et al., 2010; Elger et al., 2004; Khan et al., 2010). Further to the inter-ictal period, the time immediately following seizures (the "post-ictal" period) has been associated with many alterations to physiological function. These alterations range from psychosis (Krauss & Theodore, 2010), which has in turn been associated with out of character violence (Kanemoto et al., 1999), to disturbances to vegetative functions (Hoffmann et al., 2009; Fauser et al., 2004; Janszky et al., 2007). For many, these ongoing alterations in epilepsy patients could be argued to have a larger impact on their quality of life, given that seizures are often infrequent, and typically of short duration, unlike the otherwise invisible impact upon cognition and personality. Epilepsy is also associated with a higher mortality rate; calculated to be between 1.6 and 11.4 times the expected rate (Holst et al., 2013; Fazel et al., 2013).

### 1.1.3. *The need for new treatment*

Despite the availability of many anti-epileptic drugs, some 30% of patients remain refractory to medication, and continue to suffer seizures (Moshé et al., 2015; Laxer et al., 2014). In a subset of the patients with medically intractable epilepsy, for whom a focal onset is able to be located through imaging or electrophysiological recordings, resective surgery is an option. However, although surgical resection is underutilized (Wiebe & Jetté, 2012), it is not a guarantee of seizure freedom, with roughly 40% of patients continuing to have complex partial or generalized seizures 12 months after surgery (Wiebe et al., 2001; Moshé et al., 2015). Even with a clearly defined focal onset zone, resective surgery is not an option for all patients, since the proximity of the focus to eloquent cortex will alter the risks from case to case. In temporal lobe resections, there is also risk to memory function, with substantial postoperative memory impairments found in up to 5% of patients (Wiebe et al., 2001).

Vagal nerve stimulation has been shown to be capable of reducing seizure frequency, though in the majority of cases does not provide seizure freedom (Moshé et al., 2015). A similar lack of seizure freedom is found in most cases utilizing deep brain stimulation (Fisher et al., 2010; Moshé et al., 2015). Recent studies have shown potential for controlling seizures using closed loop, seizure detection and subsequent intervention systems (Wykes et al., 2012; Paz et al., 2012; Ehrens et al., 2015). While necessarily invasive, such techniques allow for the treatment to be delivered only when necessary, unlike pharmacological intervention, thereby avoiding secondary effects arising from treatment during inter-ictal periods. This benefit may come at a cost however; since the detection systems take time and thus are unlikely to interrupt inter-ictal bursts, leaving the associated cognitive deficits potentially unaltered. In any case, further understanding of seizure mechanisms, including ictogenesis, seizure propagation and seizure termination, will be of considerable benefit to developing new treatments that are greatly needed.

## 1.2. The electrographic profile of seizures: the ictal core and penumbral territories

Multiple mechanisms have been posited for the seeming unpredictable ictogenesis in unprovoked seizures. It is beyond the scope of this thesis to do justice to the depth and breadth of the work into seizure onset, but importantly, the mechanism is likely to differ

from patient to patient, and indeed, there is little evidence for any distinctive pre-ictal activity that may be used as a predictor (Mormann et al., 2007).

Once ictal activity is present, however, animal models have long suggested characteristic features of cortex that has been recruited to the seizure, and the territory surrounding it. These studies describe epileptiform activity in terms of a localized region in which there are intense, hypersynchronous discharges, with many neurons showing paroxysmal depolarizing shifts (PDSs) (Kandel & Spencer, 1961a; Matsumoto & Marsan, 1964; Traub & Wong, 1982). This focal region of epileptiform activity is surrounded by territory in which there is a marked inhibitory response, restraining the propagation of the focal pathophysiology (Prince & Wilder, 1967; Dichter & Spencer, 1969a; 1969b; Wong & Prince, 1990; Schwartz & Bonhoeffer, 2001; Timofeev et al., 2002; Timofeev & Steriade, 2004; Trevelyan et al., 2006; 2007; Cammarota et al., 2013). These two areas of qualitatively different patterns of activity during seizures have been termed the "ictal core" and "ictal penumbra" respectively (Schevon et al., 2012; Trevelyan & Schevon, 2013; Weiss et al., 2013; 2015).

### 1.2.1. *The penumbra: inhibition at the wavefront*

The ictal penumbra in some instances will be subsequently recruited to the ictal core; the feature that separates this region from the rest of the cortex is its marked inhibitory response that appears to oppose the propagation of the core at that moment in time. In studies of mouse brain slices *in vitro*, pyramidal cells ahead of the ictal wavefront receive intense levels of inhibition, which appear to restrain the propagation of the ictal core by suppression of the intense excitatory drive (Trevelyan et al., 2006; 2007). This penumbral inhibition is in keeping with similar activity recorded *in vivo* (Prince & Wilder, 1967; Dichter & Spencer, 1969b; Schwartz & Bonhoeffer, 2001), and this inhibitory restraint has since been found in humans undergoing invasive monitoring prior to resective surgery (Schevon et al., 2012).

It is hypothesized, therefore, that seizure propagation occurs when this inhibitory restraint fails, allowing for the downstream pyramidal cells to be recruited due to the intense excitatory drive. The exact mechanism for the failure of this inhibitory restraint is a topic of ongoing study. It is possible that there is a simple depletion of gamma-aminobutyric acid (GABA), the neurotransmitter released by cortical interneurons that results in the postsynaptic inhibition (Zhang et al., 2012). Alternatively, the parvalbumin containing fast

spiking interneurons, which are responsible for the feed forward inhibition, may enter depolarization block (Ziburkus et al., 2006; Cammarota et al., 2013), whereby the membrane potential of the cell becomes so depolarized such that voltage gated sodium ion channels fail to de-inactivate (Somjen, 2004). This leaves the cell incapable of returning to a sub-threshold membrane potential, and incapable of firing (Bragin et al., 1997).

Another potential cause for the failure of the inhibitory restraint is through accumulation of chloride ions within the postsynaptic cell (Kahle et al., 2008; Dzhala et al., 2010; Kaila et al., 2014). $GABA_A$ receptors, which are found in the postsynaptic cleft and mediate the inhibitory response through GABA, do so via allowing an influx of anions into the cell, most commonly chloride ions. Due to the negative valence of the chloride ions, this results in an inhibitory postsynaptic potential (Nicoll et al., 1990). However, the inhibitory effect in the postsynaptic cell is dependent on the chloride ion reversal potential; as the intracellular concentration of chloride ions increases, the concentration gradient decreases, resulting in less drive for chloride ions to enter the cell. This results in a reduction in the hyperpolarizing effect of GABA (Thompson & Gähwiler, 1989a). The intracellular concentration of chloride ions can rise either from altered KCC2 expression (the potassium-chloride co-transporter that actively extrudes chloride ions from the cell; Huberfeld et al., 2007), or from intense GABAergic stimulation (Thompson & Gähwiler, 1989a; 1989b; 1989c; Dzhala et al., 2010).

In extreme cases the intracellular concentration of chloride ions may become so raised such that GABAergic transmission in fact becomes depolarizing rather than hyperpolarizing (Cohen et al., 2002; Khalilov et al., 2003; Dzhala et al., 2010). Importantly, it has been shown that, in both post-tetanic stimulation and low $Mg^{2+}$ *in vitro* slice preparations, local application of GABA produces depolarization of pyramidal cells resulting in bursts of action potentials, rather than inhibitory postsynaptic potentials (Fujiwara-Tsukamoto et al., 2006). Cohen et al. (2002) found that the subset of pyramidal cells that were bursting during spontaneous inter-ictal like activity in resected tissue from patients with temporal lobe epilepsy were depolarized by GABAergic activity. Shortly after, Khalilov et al. (2003) found that the induction of an epileptogenic focus was associated with a long term alteration to GABAergic synapses, leaving them excitatory due to the shift in the chloride reversal potential.

Furthermore, Ellender et al. (2014) showed that, during ictal discharges in *in vitro* hippocampal slices, optogenetic activation of the normally inhibitory parvalbumin-

expressing interneurons using channelrhodopsin led to an increased frequency of afterdischarges in pyramidal cells. Moreover, during the same study, the authors found that optogenetic inhibition of parvalbumin-expressing interneurons through archaerhodopsin resulted in a reduced number of afterdischarges, showing that through the intracellular accumulation of chloride ions in pyramidal cells, the postsynaptic effect of GABA can become depolarizing, thereby driving activity rather than inhibiting it. This switch to a depolarizing effect of GABAergic transmission likely plays a key role in the propagation of ictal activity through cortical tissue as it becomes recruited to the ictal core.

### 1.2.2. *The boundary between the ictal core and penumbra: an issue of seizure localization*

This region of feed forward inhibition demarcates the boundary between recruited and non-recruited cortex during a seizure (Trevelyan et al., 2006). This, however, does not suggest that the penumbral territory shows no alteration in activity. Importantly, the presence of the intense synaptic activity in the penumbra will result in large amplitude activity in the lower frequencies; notably those recorded by electroencephalography (EEG) and electrocorticography (ECoG). This high level of low frequency activity, however, is not accompanied by concomitant firing of the postsynaptic cells in the penumbra (Schevon et al., 2012; Merricks et al., 2015). The frequencies most closely related with neuronal action potentials are greater than 300 Hz, and thus not readily assessed through ECoG due to spatial averaging, and often, sampling frequencies used. Moreover, in the use of standard EEG this is further compounded by the attenuation of such high frequency activity by the skull.

As a result, standard procedures for localizing the seizure onset zone prior to resective surgery are incapable of distinguishing between cortex that is truly recruited to the pathological activity, and that which is preventing the propagation of the seizure, and is presumably, healthy. An analysis of micro-electrode array (MEA) recordings in humans has indeed found this discrepancy between the two regions, with large amplitude EEG but low levels of firing in the penumbra, and a sharp delineation at the boundary between the recruited and non-recruited cortex (Schevon et al., 2012).

The necessity to define clearly these two regions using more widely available, and ideally less invasive, tools is evident. Firstly, resection of the putative ictogenic cortex carries with it

substantial risks to the patient, and the removal of more tissue than necessary through the inclusion of the penumbra increases potential negative outcomes unnecessarily. Secondly, as mentioned in section 1.1.3, the number of patients in whom seizure freedom is maintained after resective surgery is low. The electrographic similarity between ictal core and penumbral territories in the low frequencies is a potential confound in determining the true seizure onset zone (Schevon et al., 2012), which is already a non-trivial task (Nadler & Spencer, 2014). This confound increases the possibility of the resection not containing the ictogenic region in its entirety. Indeed, a post hoc analysis of the outcomes of these surgeries has found a good correspondence between the extent to which markers of the ictal core were incorporated and the resultant seizure freedom (Weiss et al., 2015).

The aforementioned study utilized phase locked high gamma activity early in the seizure as a biomarker for the ictal core, since the high gamma frequencies are indicative of multi unit firing, and thus the locking of bursts of these high frequencies to the lower frequencies, which are associated with synaptic activity, is indicative of recruitment (Weiss et al., 2013). The phase locked high gamma is a significant progression in seizure localization tools, enabling assessment of territories that are recruited very early in the seizure. To tease apart the mechanisms that underlie this recruitment, and in turn to disentangle the early ictal core from the original seizure onset zone, requires further analysis of these activities.

## 1.3. Extracellular recordings of action potentials

As discussed in the opening of this chapter, action potentials are the means by which neurons send information rapidly throughout the brain. Analysis of neuronal firing patterns can therefore tell us a great deal about brain function, particularly for neurons close to the output of the motor system, or early in the sensory system (Hubel & Wiesel, 1977). It is also possible to scrutinize these action potential trains for clinically relevant information. For instance, in the context of epilepsy, rather than questioning how neuronal firing relates to the outside world, one might ask what the firing patterns tell us about the local brain state. How do they relate to the timing and frequency of both inter-ictal events and seizures (Verzeano et al., 1971; Ishijima, 1972; Babb & Crandall, 1976; Schmidt et al., 1976; Babb et al., 1987; Bower & Buckmaster, 2008), as described in sections 1.1 and 1.2?

### 1.3.1. *The local field potential*

It is unfeasible to record singularly from individual neurons non-invasively. The EEG, recorded from the scalp, results in poor spatial resolution at the scale required for single cells, in part due to the filtering properties of the structures between the cortical surface and the recording electrodes in EEG. As a result, there is little relationship between EEG traces and the underlying action potentials from neurons in the region (Nunez, 1981; Niedermeyer & da Silva, 2005). Whilst overcoming the spatial limitations of EEG to a certain extent, and also allowing for accurate recordings at higher frequencies, the equivalent electrodes laid onto the cortical surface (electrocorticography; ECoG) still lack the spatial precision to record activity from individual neurons.

With recordings from electrodes placed within the cortical tissue itself, however, we may record the activity of individual neurons in the near vicinity. These extracellular recordings are referred to as the local field potential (LFP), though importantly, at lower frequencies these traces may be far from local, with many cortical regions showing entrained activity (Buzsàki et al., 2012). That said, at higher frequencies, LFP traces from adjacent electrodes, separated by 1 mm, can show discrepancies (Destexhe et al., 1999), and more recent recordings from primary visual cortex in cats have suggested that the LFP in fact arises from within 250 µm of the electrode tip (Katzner et al., 2009), though this is, of course, dependent on the electrode impedance. Indeed, this scaling of spatial specificity with frequency has been posited as a hierarchy on which cortical processing occurs, allowing for precise local activity in the high frequencies in differing areas to have computational points of reference in the lower frequencies (Buzsàki & Watson, 2012; Buzsàki et al., 2013; Buzsàki & Draguhn, 2004).

How then, might these local recordings provide insights to when action potentials occur in specific neurons in the immediate vicinity of the electrode tip? Individual action potentials in fact have little bearing on the raw LFP, with the overall signal instead arising from a superposition of a multitude of trans-membrane currents in the area, ranging from synaptic activity, to interactions between neurons and glial cells (Buzsàki et al., 2012). Whilst the main contribution of single action potentials to the LFP is through the consequent postsynaptic effects, alterations to the LFP at the moment of an action potential are likely limited to those from neurons with their cell bodies within a 50 µm radius of the electrode tip (Henze et al., 2000), due to the filtering properties of the cortical tissue itself. Given these

limitations, in the event that spiking neurons are present within this tight radius to the electrode tip, these extracellular recordings of the LFP contain short (< 2 ms), low amplitude (< 1 mV) deflections that correspond to the flow of ions between the extracellular and intracellular space during the action potential.

To date, the majority of these acute recordings from within the cortex in humans have been performed using hybrid depth electrodes that incorporate stainless steel or tungsten microwires with the standard clinical contacts (Fried et al., 1999). At present, three neurosurgical institutions are also utilizing silicon based microelectrode arrays ("Utah array"; Waziri et al., 2009; Truccolo et al., 2011). A considerable benefit of using hybrid depth electrodes can be found in their ability to record from deeper structures than their microelectrode array counterparts, such as the amygdala and hippocampus (Rutishauser et al., 2006). Furthermore, the currently available Utah arrays for use in humans are configured with all electrodes at a consistent length, and so are limited to recording from a single layer within the cortex, whereas the aforementioned microwires are able to be trimmed to differing lengths so as to record simultaneously from multiple layers.

That said, the Utah array allows for a higher density of electrodes, and the rigid structure of these electrodes gives confidence of their ultimate location post-implantation in the cortical tissue, along with higher spatial sampling. Hybrid depth electrodes are advanced slowly, with the microwires spreading outwards from the tip, whilst the Utah array is implanted quickly with a pneumatic insertion technique (Waziri et al., 2009). The slow insertion of depth electrodes provides the benefit of being able to populate a distribution of activity from cells that are different to those that are ultimately recorded from. This distribution can be populated within electrode prior to advancing the electrode further through the tissue, thereby ensuring the distribution has arisen from different cells, and this provides a "null" distribution corresponding to the activity that may be detected purely by chance in the tissue, against which recording stability through time can be assessed. This limitation in the Utah array implantation procedure can be overcome through using one or many of the other electrodes to derive the equivalent "null" distribution, as described further in section 3.3.1. Finally, the flexibility of microwires in the hybrid depth electrodes potentially allows for a reduced tissue response to the insertion due to less shearing, but simultaneously reduces stability of recording location through time. Ultimately, both electrode systems provide the ability to record the brief spikes corresponding to action potentials in nearby neurons, and

the decision of which to use is controlled by the location of the desired recordings, and clinical considerations.

### 1.3.2. *Spike sorting*

Given the presence in the LFP of these temporally brief spikes corresponding to action potentials in nearby neurons, how can they be isolated, first from the much higher amplitude ongoing synaptic activity, and second, to individual neurons? The larger amplitude oscillations resulting from the synaptic activity are present predominantly below 500 Hz, and with much greater power in the lower frequencies (Buzsàki & Draguhn, 2004). Extracellular action potential recordings, however, are complete within 2 ms, thereby having much less of a component below 500 Hz. Indeed, the removal of all isolated spikes from the LFP has no significant effect on the power in frequencies below 100 Hz (Belluscio et al., 2012). As a result, applying a high pass filter to the LFP at roughly 300 Hz will remove the majority of synaptic activity from the signal, without having a major impact on the deflections that correspond to action potentials from the nearby cells.

The high pass filtered signal therefore corresponds to the spiking activity of multiple neurons, with larger amplitude deflections corresponding to the neurons with cell bodies near to the electrode. Further insight into the relationship between action potentials from individual neurons within a population, to behaviour and disease, may require the separation of these spikes into their respective neuronal sources. Within cell type, action potentials show similar features from cell to cell, but the extracellular spikes vary depending on the exact geometry between the neuron and the recording electrode (Gold et al., 2006).

Since two neurons cannot inhabit the same space, these alterations to the spike width, amplitude, and shape allow for separation of spikes into groups of similar waveforms, which putatively arise from individual neurons. Whilst less reliable than cell-attached or intracellular recordings, this putative clustering into separate neuronal sources allows for simultaneous recordings of multiple neurons, and importantly, these LFP recordings are feasible in humans. The simultaneous recordings from multiple, putatively separated neurons in turn allow for network spatiotemporal activity patterns to be probed in various situations, including during, and in the time points surrounding, epileptiform activity in humans as described here.

Importantly, however, these extracellular recordings prove difficult to further classify each neuron by cell type. As discussed, the complexities of cortical processing, and of seizure onset and propagation, lie within the intricate interactions of different cell types within the neuronal network. The methods and limitations of performing the putative classifications into single neuronal sources, and subsequent putative sub-classification by cell type, are discussed in depth in chapters 2 and 3.

## 1.4. The rationale of this thesis

Studies of the spatiotemporal activity patterns of multiple individual neurons from MEAs in recent years have provided great insights to methods of cortical processing, such as the ability to predict a visual stimulus from the response of individual neurons in humans (Quian Quiroga et al., 2005; 2007). These are evidently important recordings in our current stage of scientific enquiry into the neuronal substrates of cognition and consciousness, though in order to utilize fully the resultant spike sorting data, we must have first a solid knowledge of how such recordings from one time point relate to those from another.

Indeed, to maximize both the power and benefit of experiments, and justify these invasive recordings given the cost, it is important to perform multiple experiments over many days during their implantation in humans. However, to compare the results from one experimental epoch to another from a separate day, we need to know the stability of putative individual neurons from one day to another – can the response of a neuron on a given channel from one experiment be pooled with the response from the same channel on other days? Whilst this required stability has been shown to be present in chronic implantations (Donoghue et al., 2007), recordings prior to resective surgery in epilepsy patients such as analysed here, are necessarily brief, occurring within the first days to weeks of implantation. No analysis of stability of the recordings from single neurons from day to day during this time period has been performed.

Furthermore, the employment of these recordings from MEAs in the presumptive seizure onset zone in humans has demonstrated the ability to spike sort during seizures (Truccolo et al., 2011). As outlined in section 1.2, this would not be expected if human seizures demonstrate equivalent features to animal models of epileptiform activity. The PDS and hypersynchrony of the ictal core activity would be anticipated to render spike sorting

implausible, given its reliance on maintained wave shape from an individual neuron. This highlights the importance of analysing all recordings from humans in these instances with specific respect to the location of each electrode – is it within the fully recruited ictal core, or does it remain penumbral to the recruited territory? This discrepancy may highlight important differences between the two regions, and not only validate long standing animal models of epileptiform activity, but likely also clarify potential insights into seizure prediction capabilities, how spatiotemporal neuronal activity corresponds to seizure onset, propagation and termination, and into the comorbidities of epilepsy.

## 1.5. The aims of this thesis

Accordingly, the main aim of this thesis is to assess the potential for continuous recordings of single units using MEAs implanted in humans during the acute / subacute (days to weeks) phase, and how these might elucidate seizure mechanisms in humans with respect to specific cortical regions.

In order to do so, the specific aims of this thesis are as follows:

(i) To characterize the level of stability in recordings from these MEAs over many hours, up to multiple days, with respect to the high frequency deflections which correspond putatively to the extracellular trace of action potentials in nearby neuronal sources;

(ii) To assess how the activity of the isolatable neurons in the vicinity of the MEA is altered during ictal events, with specific respect to the ictal core and penumbral territories, and how this relates to long standing animal models; and

(iii) To relate the spatiotemporal activity patterns of these neurons during the post-ictal period to the ongoing clinical symptoms, thereby querying the relationship between single neurons and the potential post-ictal altered consciousness. In doing so with respect to the ictal core and penumbral territories, these features will also be assessed for their potential utility as a seizure localization tool.

# Chapter 2. Materials and Methods

## 2.1. Recording Methodology

The majority of data presented herein were collected from recordings conducted by the treating physicians and surgeons at Columbia University Medical Center/New York Presbyterian Hospital in human patients with intractable epilepsy. The recordings were performed under the oversight of Columbia University's Institutional Review Board (IRB), complying with all regulations, and informed consent was obtained from each patient. Data from these recordings were shared under the provision of the Columbia University IRB protocol, and all recordings having been de-identified. All analyses of these recordings were performed offline, in the MATLAB environment (MathWorks, Natick, Massachusetts). Boxplots throughout were created either in the MATLAB environment, or using online software (http://boxplot.tyerslab.com; Figures 4.6 & 4.12).

### 2.1.1. *Patient cohort*

Adults with pharmacoresistant focal epilepsies, with indication of a sublobar seizure onset zone from presurgical evaluation, underwent chronic invasive monitoring with standard subdural electrocorticography (ECoG) sensors to aid localization of the ictogenic region for subsequent resection in order to control seizures. Seven of these patients were simultaneously implanted with a 4 mm x 4 mm, 96-microelectrode array (MEA) designed to record from layers 4 and 5 of the lateral neocortex (Neuroport, Blackrock Microsystems Inc., Salt Lake City, Utah), for the duration of the invasive evaluation (a range of 4 to 28 days' total implantation duration per patient). Patients were deemed eligible for the simultaneous MEA implant if the presurgical evaluation indicated clear seizure localization and thus the subsequent invasive monitoring was used to further refine the resection boundaries allowing for the implant location to be within the putative seizure onset zone. The implant location was selected to be away from eloquent brain regions, so as to ensure the implant location would be within the subsequent resection.

Patients included both temporal and frontal lobe epilepsies; lateral temporal lobe sites were selected to be within the anticipated anterolateral temporal lobectomy. Implantation sites were selected based on the presurgical evaluation, including video electroencephalography (EEG) recordings, and intraoperative ECoG recording performed prior to device implantation, that indicated a well-defined inter-ictal and ictal focus (Schevon et al., 2012). Since the purpose of these studies was to assess the spatiotemporal activity patterns in the peri-ictal period, analyses were limited to 4 of the 7 patients, in whom paroxysmal seizures were recorded. Of the three patients excluded, one had no seizures during the 2 weeks of recording, one patient remained in status epilepticus throughout the 48 hours of recording, and there were technical difficulties with the MEA after the first 24 hours in the final patient. Specifically, this patient experienced a small haemorrhage under the subdural ECoG grid, which was an anticipated potential adverse effect of the surgery and not related to the implantation of the MEA. Of the four patients included, two were female and two male, ranging in age from 19 to 39 years. All four patients had complex partial seizures; patient 2 also experienced secondary generalization (see Table 2.1 for further details).

### 2.1.2. *Microelectrode array and recording details*

The MEA used for these recordings consisted of 96 microelectrodes arranged in a square 10 by 10 grid, without electrodes in the corner positions, with a 400 µm inter-electrode separation (Figure 2.1 A). Each microelectrode was made of platinum-coated silicon, with a length of 1 mm from the base and electrically insulated up until the last 70 µm from the tip. The electrodes tapered from between 35 and 75 µm at the base, to between 3 and 5 µm at the tip. Electrode impedance at manufacture was 322 ± 128 kΩ (Blackrock Microsystems, Utah). The 1 mm length electrodes were anticipated to result in the recording tips being located in layer 4/5, based on the recording locations from the “BrainGate” research team (Donoghue et al., 2007). The MEA implantations by Donoghue et al., however, were in primary motor cortices, a location with greater grey matter thickness than other cortical regions (Douglas et al., 2004). Note also, that cortical regions that are regularly recruited to ictal activity may show thinning of grey matter due to neuronal loss. With the aforementioned considerations, it was anticipated that the electrode tips would be in layer 4/5, and this was confirmed to be the case in at least one instance through post-resection immunostaining (Figure 2.1 B).

| | Patient 1 | Patient 2 | Patient 3 | Patient 4 |
|---|---|---|---|---|
| **Age/Gender** | 32/F | 19/F | 39/M | 30/M |
| **MEA location** | Left inferior temporal gyrus; 2.5 cm from anterior temporal pole | Right posterior temporal; 1 cm inferior to angular gyrus | Left lateral frontal; 2 cm superior to Broca's area | Left supplementary motor area; 3 cm superior to Broca's area |
| **ECoG location** | Left lateral & subtemporal | Right lateral and subtemporal, parietal, occipital | Left lateral and mesial frontal | Left lateral frontal, mesial frontal, temporal |
| **Seizure onset zone** | Left basal/anterior temporal | Right posterior lateral temporal | Left frontal operculum (3x3 cm cortical area) | Left supplementary motor area |
| **MEA seizure territory** | Ictal core | Ictal core | Penumbral | Penumbral |
| **Days recorded** | 5 | 28 | 4 | 4 |
| **Seizures captured** | 3 | 1 | 7 | 22 |
| **Seizure type** | Complex partial | Complex partial with secondary generalization | Complex partial | Complex partial/tonic |
| **Pathology** | Mild CA1 neuronal loss; lateral temporal nonspecific | Nonspecific | Nonspecific | N/A (multiple subpial transections performed) |
| **Outcome (> 2 years)** | Engel Ia | Engel Ia | Engel Ia | Engel III |

**Table 2.1. Patient details**

Further information about the patients may be obtained by contacting the treating physician, Dr Catherine Schevon (Department of Neurology, Columbia University Medical Center).

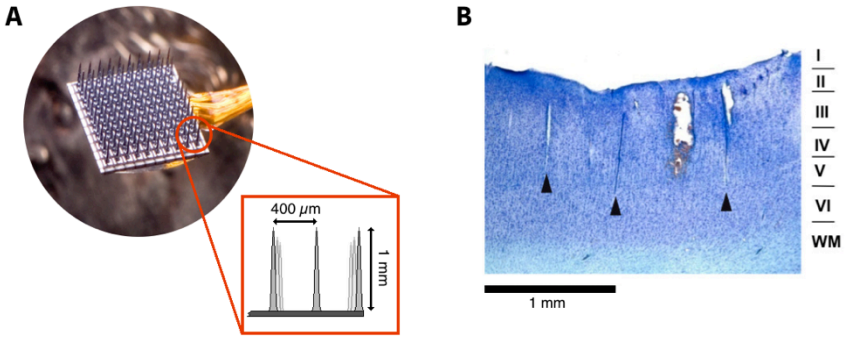


**Figure 2.1. Microelectrode array (MEA) technical details.**
**A.** A highly magnified photograph of the MEA used throughout these studies. The MEA consists of 100 electrodes, arranged in a 10 x 10 grid, with the four corner electrodes acting as ground in order to shield the remaining 96 active signal microelectrodes. Note that these corner electrodes were not used as reference, but rather two reference wire options were employed: one subdural and one epidural. Red box shows a cartoon representation of the individual electrodes, with the distances to scale, displaying the 1 mm shank, and 400 µm electrode pitch. Only the final 70 µm tip of each electrode was active, with the remainder of the electrode being electrically insulated. **B.** A Nissl stained section from patient 3's MEA implantation site. Note that the electrode tracks are clearly visible (black arrows), and these indicate that the position of the electrode tips was in layers 4/5. Parts A and B are modified and reproduced from Schevon *et al.* (2012) respectively.

The MEA was implanted through the pia mater into flat surfaces of neocortical gyri, using a pneumatic insertion technique (Waziri et al., 2009). Two reference wires were available; one was placed subdurally and the other epidurally, with the wire resulting in higher signal to noise ratio chosen as the reference.

Recordings were started after a recovery period of one day, and starting two days after implantation patients' antiepileptic medication was gradually reduced where necessary to provoke habitual seizures, one medication at a time. The reduction of medication was halted in the event of high levels of epileptic activity or the occurrence of seizures.

The timings of seizures were detected by the treating clinicians from clinical symptoms during monitoring, and further confirmed by post hoc analysis of the EEG recordings by two neurophysiologists at New York Presbyterian Hospital, Drs. Schevon and Emerson. This detection of seizure onsets and offsets was done using standard electrographic criteria.

Signals were recorded from the MEA at 30 kHz per channel with 16-bit precision, and a range of ± 8 mV. Data were band pass filtered between 0.3 Hz and 7.5 kHz. Subdural ECoG signals were recorded from standard clinical strips or grids of platinum macroelectrodes of 4.0 or 4.5 mm diameter, with 1 cm spacing between centre points. Data from the macroelectrodes were recorded at a sampling rate of 500 Hz to 2 kHz per channel, with 24-

bit precision, and a band pass filter between 0.5 and 125 Hz (XLTek Inc., Oakville, Ontario, Canada). Epidural contacts facing away from the dura mater were used as the reference for these recordings. MEA and macroelectrode recordings were synchronized through alignment of coded digital pulses to both recording systems.

## 2.2. Spike Sorting Methodology

### 2.2.1. *Spike detection*

Multi-unit activity was extracted from each channel of the MEA using a second order Butterworth filter with a band pass between 300 Hz and 3 kHz in order to remove synaptic activity from the extracellular recordings. The use of a Butterworth filter is advantageous due to its flat frequency response in the pass band, thereby not exaggerating or diminishing components of differing frequencies within the wide pass employed here. Since the Butterworth filter is causal, and therefore introduces a phase shift, the MATLAB function "filtfilt" was employed; this method filters the signal twice, once forwards and once backwards, so as to remove the phase delay introduced by the filtration. As the extracellular deflections caused by action potentials are typically complete within 2 ms, the resultant spike can be expected to be 500 Hz at its slowest. The slightly slower, 300 Hz was chosen as the high pass end of the filter so as to allow some leeway in the timings of late components in the waveform that may provide information as to the speed of return $K^+$ ion currents towards the end of the detected waveform. Furthermore, epilepsy related depolarizations may be longer in duration (Kandel & Spencer, 1961a), and so slower than standard physiological frequencies in these recordings might be expected. Whilst extracellular spikes can significantly alter the power down to below 150 Hz (Belluscio et al., 2012), incorporating signals much slower than 300 Hz would likely confound spike sorting due to the presence of high gamma rhythms at these frequencies (Buzsàki & Draguhn, 2004).

At higher frequencies, incorporating signals faster than 3 kHz adds greater background noise relative to extra information from an individual action potential. Greater than 3 kHz activity would correspond to a unipolar deflection that is over in less than one sixth of a millisecond (as a dipolar, full period would be over in less than one third of a millisecond), thereby

adding more variability between spikes originating from a single neuron than discrepancies between spikes originating from multiple cells.

An estimation of the standard deviation of the background noise in the resultant multi-unit signal, $\sigma_n$, was calculated as the median of the absolute deviation of the signal *x* divided by 0.6745 (Donoho & Johnstone, 1994; Quian Quiroga et al., 2004; note that 0.6745 times the standard deviation of a Gaussian distribution equates to the probable error, i.e. 50% of data will fall within ± 0.6745 SD), as shown below:

$$\sigma_n = median\left(\frac{|x|}{0.6745}\right)$$

This method of estimating the background noise is distorted less by the presence of large amplitude spikes in the signal in its use of the median rather than the mean, and is therefore more robust to differences in firing rates in other epochs and electrodes than using the raw standard deviation; safely assuming that the firing rate during the inter-ictal period, on which this estimation was performed, is not high enough to leave the large amplitude action potentials as a large proportion of the total data samples (Quian Quiroga et al., 2004).

Negative deflections in the signal that surpassed 4 times this estimate of the background noise were detected in order to locate the timings of action potentials from nearby cells. Negative deflections were used for this detection because the main extracellular signature arising from an action potential in a nearby cell is the flow of positively charged $Na^+$ ions into the axon initial segment and away from the electrode tip in the extracellular space. The threshold was set to 4 times the level of background noise so as to provide a balance between false positive (incorporating deflections arising purely from the Gaussian distribution of background noise, that are not cluster-able single units) and false negative detections (not including events from single units due to their not reaching the threshold for detection). In practice, this threshold strongly favours false positive detections, which is preferable to favouring false negatives, since it is inclusive, and the background multi-unit activity that cannot be ascribable to any single unit can be removed during the clustering process, as described in section 2.2.3.

An *n* x *m* matrix of waveforms was created from these timings where *n* is the total number of detections, and *m* is the number of samples stored for each waveform. A 1.6 ms epoch

surrounding each detection, from 0.6 ms prior, to 1 ms post-detection, was used; this duration allowed for full incorporation of the extracellular spike, without introducing enough of a time window either side to significantly increase the likelihood of incorporating an action potential from a separate cell within the waveform. All spikes were inspected visually for any outliers of clear artefactual origin, nominally those that surpassed an absolute voltage of 300 μV at any time during the spike.

Since spike sorting was performed offline, post hoc, detection through the location of negative peaks was computationally feasible, as opposed to detecting the moment of threshold crossing as is necessary when performing online spike detection. As a result, spikes were already aligned to their negative peak at the sample rate used. Further alignment was performed by up-sampling the matrix of waveforms to a sampling rate of 120 kHz through cubic spline interpolation, and re-aligned to the up-sampled time point with maximal negative deflection within a 0.2 ms window of the original detection. While doing so allowed for a minutely higher precision in determining the exact time of the waveform, it was primarily performed with the aim of further characterizing the exact spike shape from each unit for subsequent clustering; it proved, however, to not add any considerable level of information to enhance the sorting of units over the already high sampling frequency of 30 kHz and was thus not used after further assessment, other than for putative sub-classification of units by cell type, as discussed in section 2.3.4.

### 2.2.2. *Electrode characteristics*

All spikes from each electrode were characterized via a compression of the information in each waveform by performing singular value decomposition on the normalized matrix of waveforms, commonly referred to as principal component analysis. This method describes the common variances present in the original dataset by de-correlating the matrix with an orthogonal transformation. These common variables are referred to as the principal components of the dataset held in the matrix.

For each waveform in the original, untransformed matrix, a vector of principal component scores is calculated that corresponds to the multiple of the corresponding principal component required in order to produce the original waveform. The matrix of principal components is ordered such that the first column corresponds to the maximal variance found in the original dataset, and the following principal components are ordered in

decreasing percentage of variance contained within each principal component. As such, typically the first *p* principal components can be used to explain a large majority (> 95%) of the total variance contained within the original matrix of waveforms. This feature is used to compress the original *n* by *m* matrix of waveforms into an *n* by *p* matrix, where $p < m$, and typically *p* is not much larger than 3.

If there were a matrix of waveforms that were to originate wholly and exclusively from *q* neurons from distinct positions relative to the recording electrode, these waveforms would also show *q* typical underlying shapes due to their relative positions from the electrode tip (Buzsàki et al., 2012; Gold et al., 2006). When these shapes are distinct from one another in any particular feature, they therefore form *q* clusters in the *p* dimensional principal component space. As is explained in the following section, this is of significant use when sorting the waveforms into groups based on their putative single neuronal sources, though was first utilized in these analyses as a method of characterizing the features of spiking patterns found on each electrode of the MEA through time.

In the instance of multiple epochs as employed in chapter 3, with each epoch of 180 second duration, detected waveforms were concatenated into a single matrix per electrode, and the principal components were calculated as a group. The first 2 principal component scores for each electrode were separated back into their respective epochs and used to describe the characteristics of the waveforms of each electrode through time. In order to assess the alterations in overall distribution of positions of the principal component scores through time, the nearest neighbour in square Euclidean distance was calculated for each data point within the principal component space for that electrode and epoch. The electrode's signature at that time point was characterized both by the average of the distance to the single nearest neighbour for each data point, and by the average distance between each data point and all other data points within that epoch.

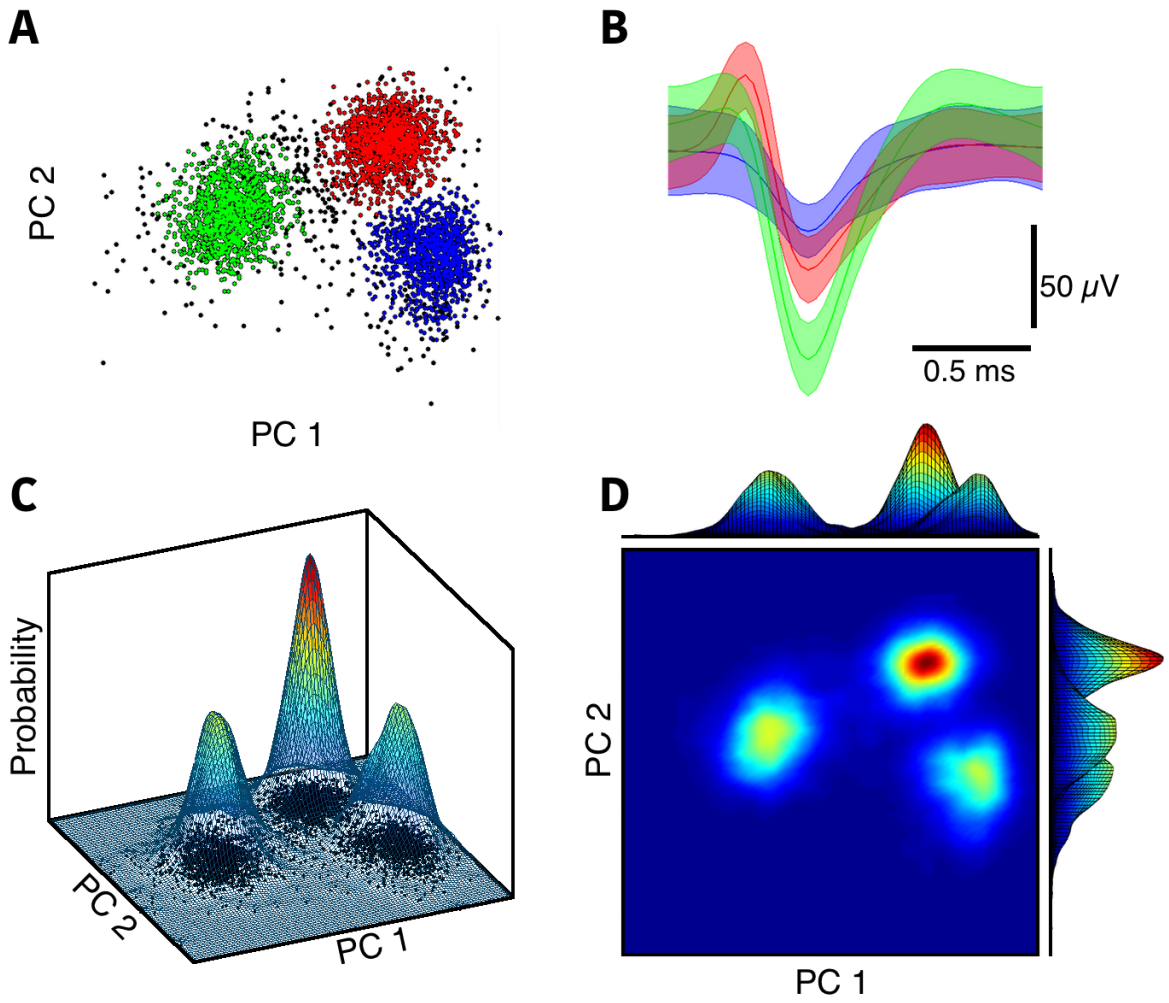


**Figure 2.2. An example analysis procedure, showing the conversion of principal component scores to probability density histograms.**
**A.** First and second principal component scores from a 3 minute epoch in patient 2. Clusters that were subsequently defined as well isolated, single units (red and green), and the background multi-unit activity that could not be further separated (blue), have been highlighted in colour. **B.** The mean ± SD waveforms from the clusters shown in principal component space in A, with colours maintained. **C.** A 3-dimensional view of the resultant histogram of the number of data points from the original principal component scores (black dots) within each bin, from a grid of 100 by 100 bins. The raw numbers in each bin is divided by the total number of data points to result in the probability of selecting a data point from that bin when choosing a data point at random. The probability density has been smoothed with a moving average of 5 bins in each direction in order to avoid the effects of isolated data points, and has been coloured from blue at the minimal probability, to red at the maximal probability. **D.** The resultant probability density plot for the principal component scores in A, with the colours for each bin maintained from C. The probabilities across the first and second principal components alone are shown at the edges.

As the results from these metrics of the nearest neighbour distribution did not correspond to the patterns seen qualitatively in the data (c.f. Figure 3.2), a separate approach was devised through calculating the probability density distribution of principal component scores through 2 dimensional space over 100 bins in each dimension, and smoothed with a moving average of 5 bins. Within electrode, every epoch's dimensions were normalized to

the maximum and minimum of the first epoch used, in each dimension in order to avoid the effects of extreme outliers in any single epoch (Figure 2.2).

In order to assess the stability of these features through time on each electrode, and the specificity of the features recorded between electrodes, two dimensional cross correlations were calculated on the smoothed probability density histograms, for both intra- and inter-electrode comparisons. The resultant maximal correlation coefficients were found to be non-normal (Kolmogorov-Smirnov test for normality; maximal $P < 0.0001$), and so the significance levels of these correlations, both through time and between electrodes, were assessed with the non-parametric Mann-Whitney *U* test in place of the standard *t*-test due to the *t*-test's dependence on the compared groups both being normal distributions. Indeed, in many instances throughout this thesis, resultant distributions have been found to be non-Gaussian, and so this assumption of normality has been avoided where possible.

A possible source of bias introduced through the use of cross correlations on the probability density histograms may arise due to particularly large levels of background multi-unit activity in the matrix of waveforms, leaving the densest probabilities being caused by a large number of stationary minor deflections that were included due to the favouring of false positives by the spike detection methods. The deflection of voltage recorded in the extracellular space from a nearby action potential, whilst influenced by the volume of the soma, is primarily dependent on the distance between the soma and the recording electrode tip, and so, the smaller the deflection, the larger the volume of the electrode's "listening sphere" from which the deflection may have originated.

As such, these deflections can be expected to arise from a significantly larger number of neurons at these lower amplitudes, with less variance in waveform shape due to the distance from the electrode tip. The result of this situation is a large, dense Gaussian distribution in principal component space showing little change over time, or between electrodes, causing a high maximal cross correlation coefficient despite possible alterations in the activity of the relatively few waveforms arising from the action potentials of nearby cells. This is especially likely to cause an exaggeration of the intra-electrode similarity despite a lack of stability in individual cells when noting that the largest deflections from action potentials recorded by these high-impedance electrodes (> 100 µV) require the source neuron's cell body to be within 50 µm of the electrode tip, with the detected voltage dropping at least an order of magnitude over 100 µm (Gold et al., 2006). This would result in

any alterations of the characteristics from any single cell that is near enough to the electrode tip to be likely lost amongst the high stability of the large number of smaller deflections arising from the background activity of many more distal cells.

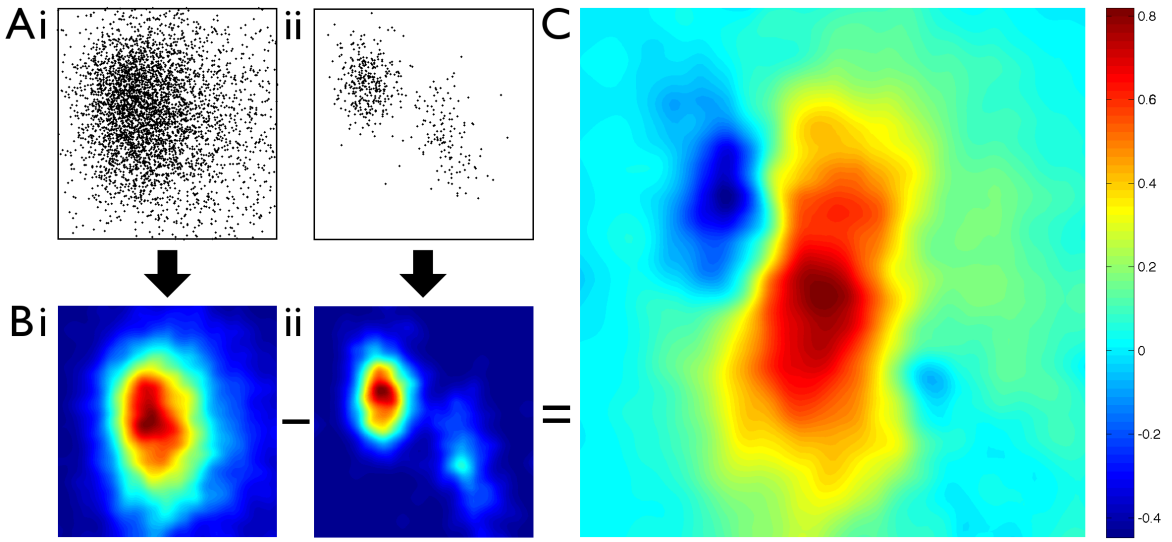


**Figure 2.3. An example analysis procedure, showing the subtractive method for the assessment of changes in unit specific activity.**
**A.** Example plots of first versus second principal component scores of spikes from an example single channel during (**i**) the seizure epoch, and (**ii**) a pre-ictal period of equivalent duration. **B.** The smoothed probability density histograms as described in Figure 2.2, of the equivalent plots above in A. **C.** The result of the subtractive method, wherein the pre-ictal epoch's probability density histogram has been subtracted from that of the seizure epoch, used to highlight changes in unit specific activity rather than simply finding correlation between the principal component scores of the multi-unit background activity from one epoch to another.

In order to avoid this biasing effect, a secondary, subtractive method of comparing the probability density histograms was devised (Figure 2.3). The probability density histogram for the first epoch of the series from each electrode was subtracted from each subsequent probability density histogram. Since the histograms were normalized in space to the first epoch's dimensions, and the number of background distal spikes that cause the large and dense, stationary Gaussian distribution was well maintained throughout the epochs, this subtractive method removes their influence from the resultant distribution. Similarly, a single neuron showing no change in either wave shape or firing rate would also result in near-zero values, leaving this subtractive method highlighting only alterations in neuronal firing rates or waveforms from cells that were close enough to the electrode tip to cause large deflections that were well separated in principal component space.

### 2.2.3. *Feature extraction*

In order to separate spikes into their putative respective neuron of origin, they were clustered based on features extracted from their spike shapes. The simplest features on which spikes may be clustered is based on aspects of the raw voltage detected during their waveform, namely the minimum, i.e. the voltage of the detected trough, and the maximum following voltage as the spike returns to baseline. While this procedure was effective on channels with a very high signal to noise ratio and large variability between few neurons, it was ineffective at separating spikes on channels with many waveforms at similar amplitudes, or waveforms with a distinctive shape that were at a similar amplitude to the background noise from more distal units with less distinctive shapes (Figure 2.4 B).

As described in section 2.2.2, the information latent in the matrix of waveforms can be compressed into fewer dimensions via principal component analysis; the number of dimensions selected for the clustering procedure is determined by the number of principal components that together explain greater than 95% of the variance in the original dataset. Using these principal component scores allows for easier separation of the waveforms that could not be separated using other simple metrics, such as spike minimum and maximum (Figure 2.4 D). The principal component scores allow for this easier separation since it takes into account every value from the detected waveform, albeit compressed in a non-lossless manner, rather than clustering solely on 2 data points whilst ignoring the other 46 from the full 1.6 ms window surrounding the spike.

Further to these features calculated by principal component analysis, another, now standard, feature compression method was simultaneously tested: the coefficients from a wavelet transform, as per Quian Quiroga et al. (2004). This method involves transforming and scaling a defined wavelet in a set number of ways, and correlating the original waveform with each of these transformed wavelets in turn, giving rise to a series of coefficients which describe the waveform as a function of each wavelet. This has been shown beneficial to clustering procedures in test datasets since it incorporates high temporal resolution, and does not rely on the features of maximal variance in the waveforms, unlike principal component analysis (Quian Quiroga et al., 2004). The compression of the dataset performed during the principal component analysis occurs through the removal of features that account for little variation in the data, and in some instances it may be these low variant features that allow for the separation between waveforms from different cells.

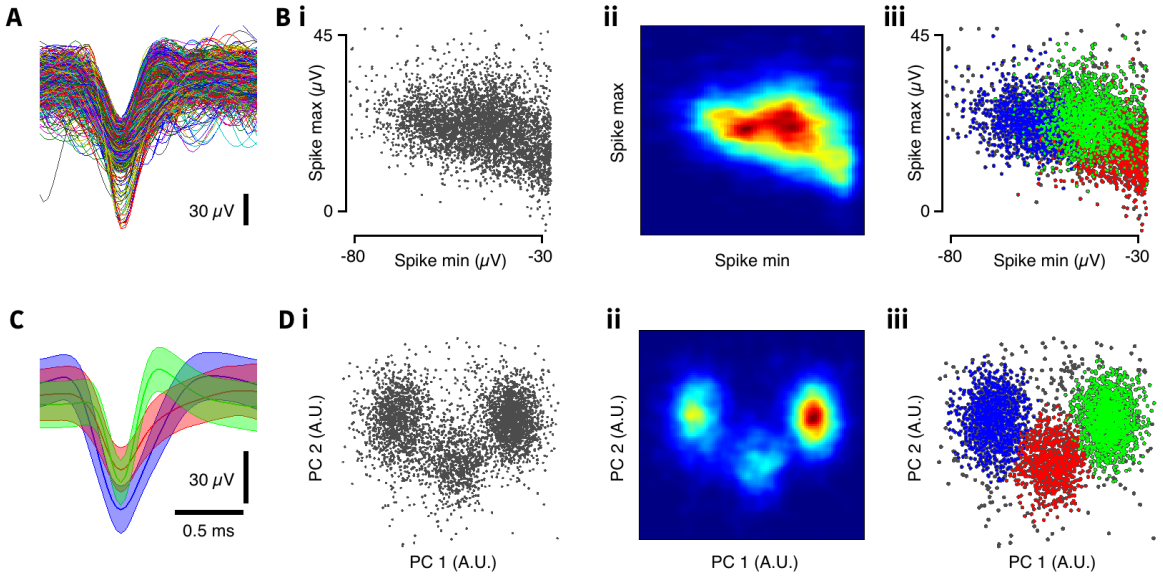


**Figure 2.4. Improved isolation of single units using principal component scores over spike minimum and maximum voltages.**
**A.** All detected waveforms from a 2 minute epoch from channel 79 in patient 2. **B.** (**i**) Spike minimum voltage (voltage at detection) plotted against spike maximum voltage for all waveforms in A. (**ii**) The probability density for the spike minimum versus maximum seen in (i). Note how no clearly isolatable clusters can be seen in either the raw minimum versus maximum, or their densities. (**iii**) Spike minimum versus maximum voltage as in (i), though colour coded to subsequently defined single units on their principal component scores (see C & D). Note that, whilst the values for each unit may be relatively well constrained within clusters, no separation between those clusters may be found, and clustering on the spike minimum and maximum in this instance would not have separated the data into these clusters. **C.** Mean ± SD from 2 well isolated units (blue and green), and the background multi-unit activity that could not be further separated (red). Units were separated on the clusters in principal component space, as seen in D. **D.** The same format as for B, though plotting the first versus second principal component scores from all waveforms in A, rather than the spike minimum versus maximum voltages. Note that clusters are qualitatively visible in both the raw principal component data (**i**), and in the subsequent density plot (**ii**). The principal component scores were clustered automatically with a *k*-means algorithm robust to outliers into 3 clusters, as described in section 2.2.4, as seen in (**iii**). The cluster colours are maintained throughout B (iii), C, & D (iii).

In the analyses used during these studies, the MATLAB code as supplied by Quian Quiroga et al. (2004) under the name “WaveClus” was used. This includes using the default parameters of the Haar wavelet, and selecting the coefficients on which to cluster as those that displayed evidence of having a non-normal distribution, as a normal distribution would be indicative of a lack of separable features from different neurons. Also in keeping with the default settings used in “WaveClus”, the 10 coefficients with the least normal distributions were selected for clustering, as calculated by the Lilliefors test, a modification of the Kolmogorov-Smirnov test for normality that does not require prior assumptions on the distribution (Quian Quiroga et al., 2004).

In some instances, a multi-dimensional feature space was compiled from both the principal component scores and the wavelet coefficients to assess whether a lack of separable clusters in one set of features could be overcome through the use of the other. This proved to add little extra useful information for clustering relative to the computational expense and so, unless specified otherwise, all clustering was performed using the first *p* principal component scores that, when combined, explained greater than 95% of the variance in the waveforms. As described previously, the extra information provided beyond these first principal component scores would be beneficial in instances where the features that could separate waveforms into distinct clusters occurred in principal components that accounted for little variance in the dataset. As such, the lack of extra useful information found in these data suggested that the optimum features for clustering were found within the first 95% of the variance. This is a situation that is likely in single units that are reliably separated; requiring information from the least varying features in the waveforms will, by definition, be relying on very minor discrepancies in waveform on which to separate into putative separate units.

### 2.2.4. *Clustering procedure*

In order to classify which waveforms belonged to which putative single unit, a *k*-means algorithm that was robust to outliers was used to cluster the features chosen to describe the waveforms, as severe outliers can heavily distort the outcome of the *k*-means algorithm. The removal of outliers in this manner was necessary, as on occasion the recordings included transient, extreme deviations that were highly unlikely to be of biological origin. Unless these are removed from the analyses, these may badly distort the *k*-means clustering results.

Using routines from the Chronux toolbox (http://chronux.org; Bokil et al., 2010), the outliers were detected by clustering the data with the standard *k*-means algorithm, selecting the data points that were significantly distal to their assigned cluster's centroid, temporarily removing those waveforms from the dataset, and then repeating a total of 3 times on the resultant cleaned dataset from each repeat (Bokil et al., 2010). The threshold for the removal of a waveform due to its distance from its corresponding centroid was set as the value at which its Mahalanobis distance from the centroid had a probability of appearing in the $\chi^2$ probability distribution of that cluster as less than $0.5/n$, where $n$ is the total number

of spikes in the cluster. That is to say, it would have appeared less than half a time by chance in the calculated distribution.

Once the non-physiological outliers were removed and stored separately for re-integration into the dataset for manual inspection later, the methods as described by Hill et al. (2011) were employed using modified scripts from the "UltraMegaSort2000" toolbox for MATLAB based on Chronux routines. The cleaned set of features were at first over-clustered, using the *k*-means algorithm to group the features into *k* clusters, where *k* is a high enough value to reasonably presume that whilst waveforms from a single neuron would be split amongst multiple clusters, no one cluster would contain waveforms from more than one neuron. In doing so, the dataset was thus split into many mini-clusters.

The *k*-means algorithm is an iterative expectation-maximization algorithm, whereby the *n*-dimensional data were assigned to one of *k* clusters. *K* data points were chosen at random from the dataset at first, to be treated as the centre, or "centroid", of that cluster. Each data point was then assigned to the cluster with the nearest centroid in the *n* dimensions in turn, with the centroid of that cluster being re-calculated after each new assignment, until all data points had been assigned to one of the *k* clusters. This process was then repeated until no data points were reassigned based on the updated centroids. Clusters that contained only a single data point were disallowed, and were assigned instead to the next cluster with the closest centroid in the feature space, resulting in potentially fewer than *k* clusters in the output.

For the initial over-clustering *k* was chosen based on a target number of spikes per cluster of 500 regardless of total number of spikes in the dataset, rounded to the closest power of 2, and limited to be within the range of 16 to 128. The lower limit ensured that too few spikes in a dataset would not cause clusters to lower the probability of their members being entirely a subset of a single neuron's spikes, with the upper limit set to ensure computational time was reasonable. In practice, the upper limit was never reached in these analyses, as a dataset of 90,510 spikes during a single clustering procedure from a single electrode would have been required to cause greater than 128 clusters being necessary in order to contain on average 500 spikes per cluster, given that the number was rounded to the closest power of 2. Even without rounding, greater than 64,000 spikes would have been required.

The stipulation that the number of clusters be a power of 2 was chosen as per the algorithm in Hill et al. (2011) so as to allow for the use of a speed optimized version of the *k*-means algorithm whereby the dataset was always split at first into 2 clusters, and then those clusters were split in 2 again to give 4 clusters, and this process was repeated until the desired number of clusters was achieved, i.e. $\log_2(k)$ times.

With the waveforms now over-clustered, the clusters were then automatically joined together if they showed weak separation from one another according to a measure of their interface energy (Fee et al., 1996). In short, as clusters arising from a single neuron are expected to form a Gaussian distribution in feature space, clusters that show a high number of data points at the interface between themselves and a neighbouring cluster would be indicative of a single cluster that had been split in two by the over-clustering procedure. This measure of similarity at the boundary between two clusters, their interface energy, was calculated as the sum of the exponent of all pairwise distances between points from the first cluster to the second, divided by a set value. The set value was calculated as one-tenth the square root of the sum of the diagonals of the within-cluster covariance matrix. These values were corrected to ignore the interface energies from a waveform to itself, all as per Fee et al. (1996) and Hill et al. (2011). This amounts to having calculated the pairwise Euclidean distances between points from each cluster, adjusted to an exponential decay, which was used to estimate the local density of the cluster. Clusters were then hierarchically recombined based on their interface energy values, until no two clusters displayed an overlap that corresponded to greater than 0.05% its own density, chosen purely as a starting point at which to begin manual inspection.

All clusters were then inspected visually and corrected for outliers of waveforms, clusters that qualitatively appeared to be incorrectly merged during the interface energy procedure or those that failed to have been combined due to the aggregation cut off despite being from the same putative neuron. Clusters that appeared to be either multi-unit activity without any separable features, those arising from the background noise of more distal cells, or any with a non-physiological mean waveform were marked to be discarded during the post hoc measurements of quality, as described in section 2.2.5.

### 2.2.5. *Post hoc accuracy metrics*

After the clusters were deemed to visually correspond to waveforms from putative individual neurons, referred to as single units since their identity as single neurons is purely probabilistic, post hoc quality metrics were employed as described by Hill et al. (2011) throughout. Only clusters satisfying all of the following criteria were accepted: (i) clean separation from all other units and the background activity of distal cells in the projection of their features onto Fisher's linear discriminant; (ii) less than 1% contamination of the absolute refractory period of 2 ms in the inter-spike interval distribution; (iii) no clear outliers of waveform shape using a $\chi^2$ probability distribution; and (iv) less than 1% of total spikes from that unit missing, as estimated by the percentage by which a Gaussian fit to the detected negative peak voltages was below threshold for detection.

The separation between two clusters was assessed visually by taking the Fisher linear discriminant of their data points in feature space, which is the projection that results in the most separation between two multivariate normal distributions (Fisher, 1936), as would be expected from clusters of waveforms from single neurons, based on the mean and covariance matrices of the two clusters. Inspecting the histograms of the data points from two clusters on this projection allowed for clusters that were from one normal distribution or not convincingly separated to be easily noticed and either manually adjusted or removed from the dataset if necessary (Pouzat et al., 2002) (Figure 2.5 A, note the clear separation between the well isolated units in (i), and the lack of separation between the two clusters that would likely arise from one unit in (ii)). It is worth noting that a lack of separation on this projection may not necessarily be indicative of lack of separation in a higher dimension in the feature space, and so further inspection was performed on plots of all dimensions in turn in cases where there was a lack of separation on the Fisher linear discriminant prior to discarding.

An estimation of the false negative numbers from each cluster, i.e. the number of spikes from the cluster that were missed during the spike detection phase, was calculated via fitting a Gaussian curve to the distribution of voltages at detection from the spikes within that cluster. As the variability in this detected voltage was expected to arise primarily from fluctuation in the underlying signal due to the Gaussian background noise, the distribution of these detected voltages was also expected to fit a normal distribution. The total estimated

false negative for that cluster was calculated as the percentage of the Gaussian curve that was sub detection threshold (Figure 2.5 B).

Each cluster was assessed for outliers in spike shape via a histogram of the Mahalanobis distances for each data point in the feature space from that cluster to its centroid, fitted with a $\chi^2$ distribution as the clusters were expected to follow a normal distribution. Any waveforms that lay considerably into the tail of the $\chi^2$ distribution were checked visually and removed from the cluster when not convincingly similar to the mean waveform from that cluster (Figure 2.5 C).

At the same time, the standard deviation on the time dimension of the matrix of waveforms from a cluster was inspected, normalized to the standard deviation of the original multi-unit signal. Any clusters that showed a much larger variability within their waveforms than that of the background noise were further inspected as this would be indicative of multiple neurons contributing to the single cluster because the level of variability during a waveform is expected to be equivalent to that of the background noise (Pouzat et al., 2002; Fee et al., 1996; Figure 2.5 D).

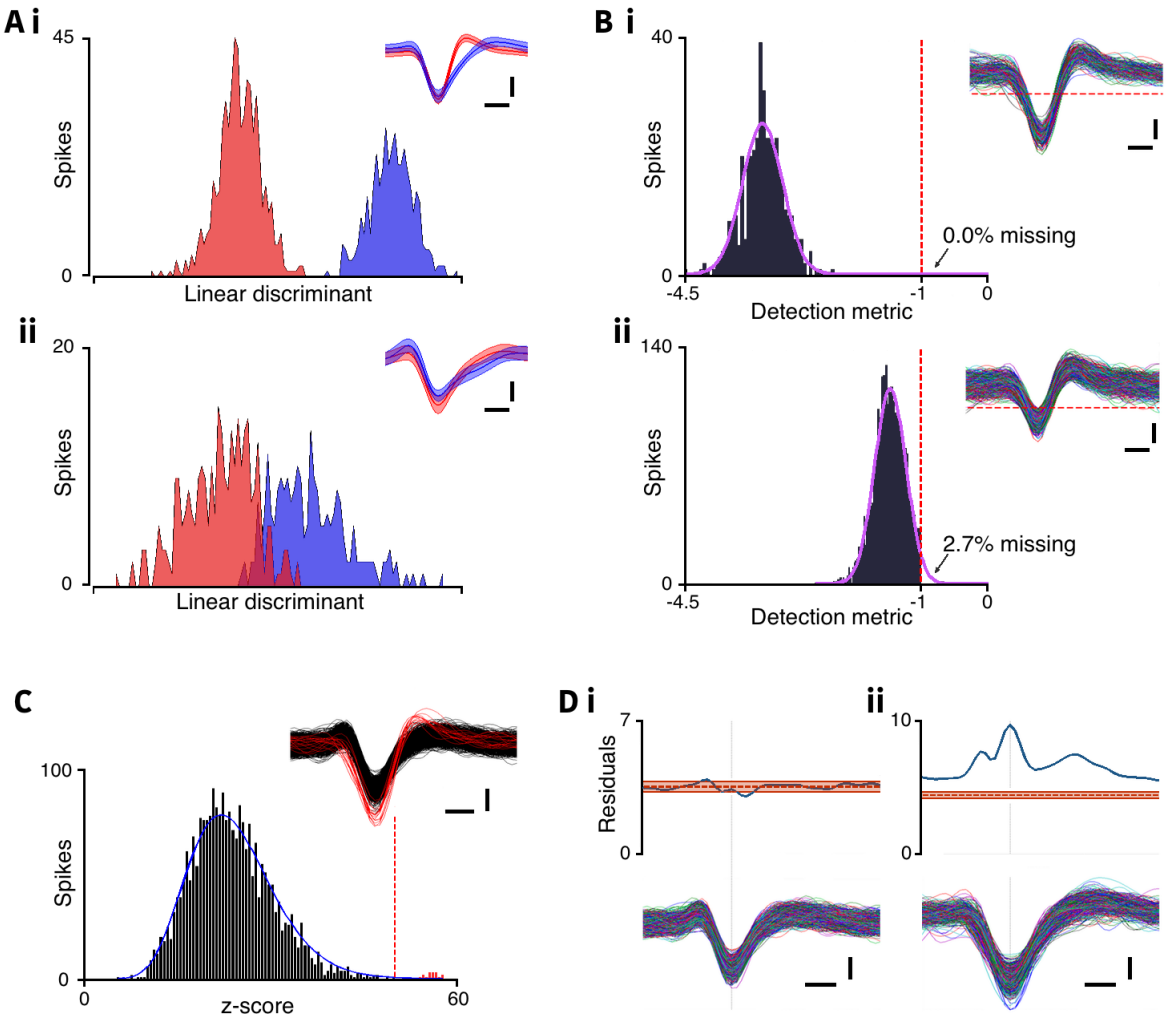


**Figure 2.5. Example results from the single unit metrics for assessing quality of clustering.**
**A.** Projection of clusters in principal component space onto the Fisher linear discriminant, showing (**i**) two units from the same channel that are well isolated, and (**ii**) a unit that has likely been incorrectly split into two, with mean ± SD waveforms inset. Note that the red and blue histograms in (ii) would form a single Gaussian distribution if combined, and the waveforms have been separated on minor differences in spike shape (inset). **B.** Measure of estimated number of false negatives from a unit. Histograms of the voltage at the detected trough of each waveform in the cluster, divided by the threshold for detection, with a Gaussian curve fitted (purple line). The voltages at detection follow a normal distribution, and so the number of sub-threshold spikes that were missed from a cluster can be estimated by the percentage of the Gaussian fit that sits above threshold. Example of a unit with a 0% estimation of false negatives is shown in (**i**), and a unit with a high likelihood of missing spikes, with 2.7% of the Gaussian curve sub-threshold in (**ii**). Raw waveforms from the respective units are shown inset, with the threshold (red, dashed line). **C.** Histogram of the Mahalanobis distance for each waveform principal component score from that unit's centroid. For a normal distribution, the Mahalanobis distances form a chi-squared distribution (blue line). Waveforms in the tail end of the chi-squared distribution are likely outliers. All waveforms can be seen (inset), with those from the left of the red dashed line on the histogram in black, and those to the right in red. Note that the red waveforms, from the tail of the chi-squared distribution, are noticeably of a different shape, and have likely been incorrectly incorporated into this cluster. **D.** Standard deviation through time (deep blue line, top) for all waveforms from a given unit (bottom). The SD of the full signal during that epoch is shown (red, dashed line), along with the 95% confidence intervals derived from the chi-squared distribution, given the number of spikes within the unit (red shading). (**i**) Example unit with no increase in variance over the background noise of the signal, and (**ii**) an example of a unit that shows more variability than would be expected without contamination from other units. Scale bars: 20 µV and 0.2 ms throughout. See Hill *et al.* (2011) for an overview of these metrics.

## 2.3. Single Unit and Local Field Potential Analyses

### 2.3.1. *Long term stability of single units*

In order to assess the ability to follow single units over long periods of time, 180 second epochs were selected every 6 hours for each channel, over a total period of 40 hours. This allowed for automated handling of the large data sizes, as storing data at 30 kHz with 16-bit precision translates to just under 10 GB of data per 30 minutes' recording time (1,800 seconds x 30,000 samples per second x 16 bit resolution x 96 channels / 8 bits per byte = 10,368,000,000 bytes ≈ 9.66 GB, without overhead or compression). For these analyses seizure epochs were excluded.

Features from the waveforms from all epochs were calculated as one in order to allow for equivalent principal components to be used throughout, as the principal components of the whole data set rather than individual epochs. Where drifting of features occurred between epochs, plots of features were inspected instead every 30 minutes over the 40-hour period to assess where altered waveforms corresponded to equivalent units (Figure 2.6). Otherwise, waveforms from all epochs were clustered and inspected simultaneously as one non-contiguous epoch.

Autocorrelation histograms over 100 ms using 5 ms bins were used to compare the typical firing patterns within units over time, as alterations in the core pattern of activity from a single unit may be indicative of their not being equivalent to the same neuron in different epochs. Autocorrelation results were used as opposed to inter-spike intervals as these distributions count when every spike from the unit occurs relative to one another as opposed to just the relationship of one spike to its precursor, thus elucidating any specific patterns found in the overall firing of the neuron; as a result, the autocorrelation will tend towards the mean firing rate of the cell, whereas the inter-spike interval will tend towards zero over longer time lags.

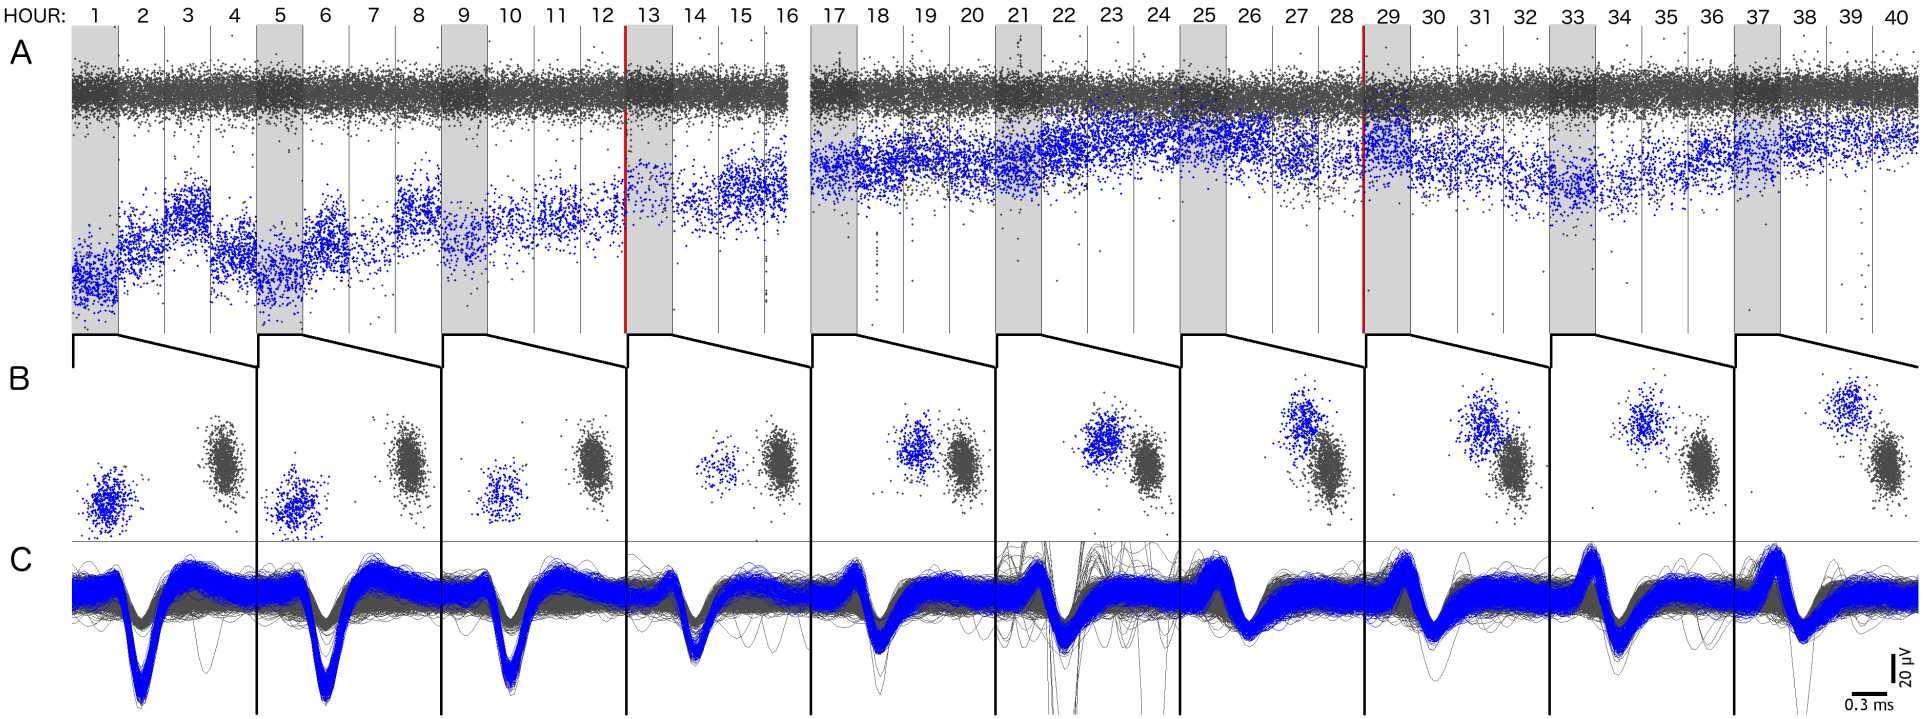


**Figure 2.6. Ability to follow drifting clusters in principal component space over a prolonged period.**
An example putative single unit that showed a large alteration in waveform over a 40 hour period, but could be traced through time as a cluster in principal component space. **A.** First principal component scores versus time, from 180 second epochs taken every hour, with background noise from distal cells in grey, and a well isolated single unit in blue (colours maintained throughout). A seizure occurred between hours 12 and 13, and a second seizure between hours 28 and 29 (red lines). The highlighted (grey background) epochs (i.e. every 4 hours) are shown in more detail in B and C. Note that the recording was interrupted during hour 16, and the pause of activity is not physiological. **B.** First versus second principal component scores for epochs highlighted in A, showing the maintained separation throughout. **C.** All waveforms from the epochs as highlighted in A, showing the change in extracellular action potential shape. Note the large alterations in waveform and associated principal component scores, but the maintained separation from the background activity throughout – whilst hours 25 and 29 show the unit approaching the background noise considerably, it can be seen in B that when including more dimensions of principal component space there is little overlap.

### 2.3.2. *Peri-ictal stability of single units*

Assessment of the peri-ictal period was performed on a 45 minute epoch from 30 minutes prior to seizure onset, until 15 minutes post onset. Timing of seizure onset was determined by the treating epileptologists, as described by Weiss et al. (2015). Waveforms from the entire epoch were clustered as a group per electrode, though in instances where the MEA was recording from the ictal core, the seizure epoch (between 30 and 120 s; patients 1 and 2) was blanked once loss of cluster-able features was confirmed, so as to allow for continuous sorting of waveforms from both the pre and post-ictal periods using the same feature space throughout, without the obstruction of the ictal core waveforms. Principal component scores, wavelet feature coefficients and spike minimum versus maximum were all checked for potential clusters after the territory under the MEA had been recruited to the ictal core, prior to blanking the seizure epoch (Figure 4.4).

To check that the loss of unit clusters during the seizure epoch in the ictal core patients was not purely obstruction in the feature space due to the increased number of waveforms, and the increased variability within those waveforms, a secondary spike sorting method was employed – template matching. Template waveforms from single units and their associated statistics were designed from the mean waveforms from units clustered in the pre-ictal 30 minutes. Every waveform detected during the seizure epoch was then compared to the template waveforms from that electrode, via both a correlation coefficient and an absolute difference between waveforms, i.e. a pairwise distance between the waveform and the template. If these values were within the range found in the same analyses from all waveforms from the pre-ictal period within the template then the waveform was assigned to that unit for inspection. After visual inspection, raw numbers of waveforms that were categorized as a single unit were calculated as a proportion of the total waveforms detected, through time, and alterations to the intrinsic properties of the waveforms were assessed, as discussed in chapter 4.

Further assessments of the waveform alterations and total numbers of spikes potentially classified as single units in the ictal core were performed by an analysis of the waveforms that were found within the boundaries of the feature space of the pre and post-ictal clusters. Cluster boundaries were defined as the convex hull surrounding the features in the 3 principal component dimensions containing the most variance explained. The first 3 were used as opposed to the total required to explain 95% of the variance due to the limitation of

fitting a convex hull in higher dimensions. This method ensured any potential complex shapes in the clusters were maintained to avoid overlap between the statistical descriptions of two clusters, and avoided potential alterations in spike shape due to including a greater number of spikes from the edges of the distribution of the cluster. This method did not allow for any alteration of the unit specific principal component scores, though in this instance this was desirable, as the template matching was used purely as a method to assess if any units maintained their intrinsic extracellular waveform whilst in the ictal core.

The full width of the spike at half the voltage of its maximal deflection (FWHM) and the spike amplitudes were calculated for waveforms that fell within the defined boundaries from the pre-ictal 30 minutes through the seizure. To check that any alterations found using these methods were not purely a result of the methods themselves, the equivalent analyses were performed on the pre-ictal time point also, as discussed in section 4.2.4.

Significant changes in the firing rate during the seizure in patients in whom the MEA remained penumbral to the seizure core was calculated by comparing equivalent duration epochs from the pre-ictal and ictal periods. Changes in firing rate were deemed significant when they fell beyond 3 standard deviations of the cell specific Poisson distribution, estimated as the square root of the firing rate divided by epoch duration (Figure 4.14; Vajda et al., 2008).

The maintenance of single units from the pre-ictal period after seizure termination was further confirmed by clustering equal duration (10 min) epochs from the two time periods independently. This was done so as to avoid the potential biasing effect of assigning post-ictal spikes to the most similar pre-ictal cluster, which may occur through a simultaneous clustering of both time periods. Once the two time periods were clustered separately, units from the post-ictal period were compared to those from the pre-ictal period. An assessment of drift in principal component space was calculated from the centroids of the two clusters. The drift coefficient corresponds to the distance between the two time points' cluster centroids, divided by the distance of the first time points' cluster centroid from the mean value in principal component space. In order to assess the significance of the drift coefficients, the true results were compared to the mean results obtained by comparing clusters to non-equivalent clusters in the other time point at random over 10,000 iterations. Statistical comparisons were made by comparing each unit to a different unit, chosen at random from the population 10,000 times, and the mean result of these comparisons was

taken as the single trial value for that unit. This strategy had the advantage of maintaining the same number of sampled units in both the measured values and the bootstrap estimates.

Further to the drift coefficient, maximal cross correlation coefficients were calculated between the mean waveform of the equivalent pre and post-ictal clusters, which was used to confirm that the two clusters did indeed correspond to the same putative single neuronal source. As a high maximal cross correlation coefficient would be expected due to all waveforms already having been selected and located based on having a noticeable negative deflection, these results were similarly compared to the mean results of comparing mean waveforms between non-equivalent clusters in the alternate time point at random, over 10,000 iterations, using the same method described for the drift coefficient comparisons.

### 2.3.3. *Post-ictal activity pattern analyses*

For analyses on the timing for recovery of single unit waveforms and alterations in typical firing patterns in the post-ictal period, a 2-hour continuous epoch from the immediate time of seizure termination was clustered along with the pre-ictal 30 minutes, where data were available. Time of seizure termination for these purposes was determined as immediately after the last large amplitude spike, with a following period of quiescence found in the high pass filtered (300 Hz to 3 kHz) signal. Activity patterns were contrasted with the subdural ECoG recordings from nearby and more distal macroelectrodes.

The immediate post-ictal period was characterized by complete inactivity of single units. The recovery of firing from single units was defined as the time of the first spike in a given cluster relative to the seizure termination. An important issue was whether this quiescence was real, or if instead it was because the unit's spike waveform had changed shape, resulting in a "classification failure" and a false impression of unit quiescence. To distinguish between these two possibilities, the principal component scores and clusters were assessed visually over time until the recovery of unit firing, to check for a drift in return of the cluster or instead a sudden reappearance of the cluster in the same principal component space as beforehand. The distribution in the delays until first classified spike within the clusters was fit with an exponential curve in order to find the time constant of the return delay of unit firing.

Coincident firing patterns between multiple units was calculated during the post-ictal period and contrasted with those of the pre-ictal period. This was done through convolving a Gaussian kernel of 1 second duration with the spike times from each unit in turn, to provide an estimation of the instantaneous firing rate of that unit through time. Differing durations for this Gaussian kernel were assessed, to ensure these features were robust to the window size. Correlation coefficients between the resultant firing rates were calculated for every unique pair of units, allowing for an assessment of the synchronicity across the population in the different time periods (Figures 5.2 – 5.5).

Autocorrelation distributions from the timings of single unit spikes were calculated through time in the post-ictal epoch and contrasted with the equivalent distributions in the pre-ictal period, in order to assess core alterations in the firing patterns of single units without reference to the underlying activity of the whole network. The autocorrelations were performed as complete 5 minute epochs in both the pre and post-ictal periods. All autocorrelations were calculated using 10 ms bins in the ± 5 s range.

Cross correlation distributions between the spike times were used to assess alterations in firing patterns at the network rather than individual unit level. These were calculated for every possible unique cell pair in the same manner as the autocorrelations. The individual cross correlations were pooled in order to gain insight into alterations in the dynamics of the network.

The mutual information present in the spike timings from all single units was calculated through time on the full epochs, using 200 millisecond resolution for mutual information calculations, focusing at zero time lag between spikes, and looking at 30 seconds of data at a time, sliding the window ten seconds at a time. All time scales used were varied to ensure reproducibility of the results regardless of the settings used. The mutual information analysis essentially examines how much does knowing the underlying pattern of activity in one series of events provide knowledge of the underlying pattern of activity in another series, giving results in bits when using log base 2. The mutual information between all unit pairs was calculated, and the mean mutual information of these results was used as the population mutual information across the network.

Bursts of activity were detected in the post-ictal period by binning the firing rate every 50 ms on 200 ms of data per window, and finding peaks in the resultant firing rate histogram that

surpassed the mean firing rate + 3 SD. Peaks that occurred within the time frame of another peak without the firing rate dropping below the mean firing rate + 1 SD were treated as one burst, and the centre of the burst was adjusted to the mean of the total detected peaks that did not lower their firing rate between. Windows of spike times of 5 s width were selected surrounding the times of these bursts, of equal duration prior to and post detection, as bursts could last up to a matter of seconds, and were often distinct on the order of seconds also, resulting in 5 s windows capturing the majority of bursts without bleed-through of other events at the edges.

Gaussian curves were fitted to the spike times within the burst windows in order to quantify the burst duration with the standard deviation of the fit, and the intensity of the burst with the amplitude of the fit. The timings of bursts were further used to create burst-triggered average activity from the nearby subdural macroelectrodes, from the wideband ECoG data in order to create a mean of power through time at different frequencies, and from the filtered ECoG data in order to retrieve mean activity in multiple frequency bands, including high gamma (100 to 250 Hz), gamma (30 to 100 Hz), beta (15 to 30 Hz), alpha (8 to 15 Hz), theta (5 to 8 Hz), delta (1 to 4 Hz), and ultra slow oscillations (0.3 to 1 Hz), so as to gain insight to the underlying activity from a broader area and from synaptic up to multi unit activity that was concomitant to the post-ictal bursts recorded in the MEA.

### 2.3.4. *Putative sub classification of single units*

Features of each single unit cluster were calculated in order to attempt a putative sub-classification of the population into fast spiking, non-bursting cells (putative inhibitory interneurons) and regular spiking bursting cells (putative excitatory pyramidal cells). A mean wideband waveform of duration 5 ms was derived for each single unit, by retrieving the spike-triggered average from the signal from the electrode on which it was recorded, filtered between 1 Hz and 5 kHz, from the spike times from the cluster. The use of a wideband waveform was used to ensure cell specific features such as the speed of the opening of $K^+$ channels after the $Na^+$ channels during an action potential, as a faster return current would be indicative of a fast spiking interneuron. These mean wideband waveforms were up-sampled from a 30 kHz sample rate to a 120 kHz sample rate through cubic spline interpolation, to allow for accuracy in the following measurements.

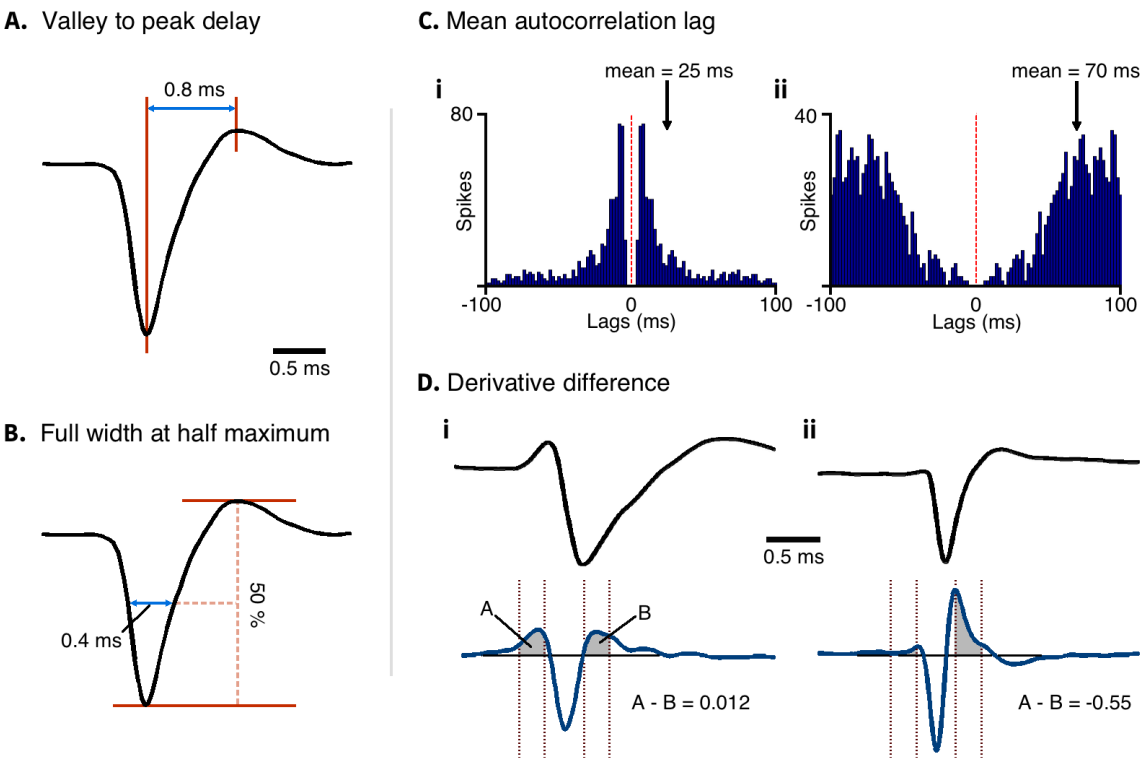


**Figure 2.7. Single unit metrics for putative sub-classification by cell type.**
**A**. The mean wideband (1 Hz to 5 kHz) waveform from a single unit (black), showing the time points (red lines) used to measure the valley to peak delay (blue, 0.8 ms in this example). **B.** The same waveform used in A, but showing the measurement of the full width at half maximum (FWHM), by taking the width of the spike halfway between its maximal and minimal voltages. **C.** Example single unit spike time autocorrelation histograms from a putative pyramidal cell due to its tendency for bursts of successive action potentials (**i**), and a putative fast spiking interneuron (**ii**). The mean values for these autocorrelations over the 0 to 100 ms range (black arrows) can be seen to differentiate these two distinctive cell-intrinsic firing patterns. **D.** The mean wideband (1 Hz to 5 kHz) waveforms (top) from a putative pyramidal cell (**i**), and from a putative fast spiking interneuron (**ii**). The high pass filtered (800 Hz to 5 kHz) first derivative show distinctive differences in the early and late components (deep blue, bottom). The difference between areas 0.25 ms in duration from ± 0.18 ms either side of the negative peak of the first derivative (grey shaded regions) were used to quantify these values, as per Csicsvari *et al.* (1998). Note that the negative peak of the derivative was used, in order to assess the spike kinetics at the specified times either side of the moment of maximal speed in the deflection of the spike, rather than once the $Na^+$ channels had already closed.

From the wideband mean waveform for each unit, multiple spike shape characteristics were derived: (i) the valley to peak duration; (ii) the FWHM; (iii) the mean autocorrelation lag; and (iv) the relative speed of onset versus offset of action potential, calculated from the area under the high pass filtered (800 Hz to 5 kHz) first derivative of the mean waveform (Csicsvari et al., 1998) (Figure 2.7 A – D respectively). These waveform features were supplemented with characteristics from the typical firing patterns of the units, comprising of the mean autocorrelation lag, as a quantification of how likely the cell was to fire in quick bursts of action potentials but less frequently outside the bursts, a feature associated with pyramidal cells, or with regularity but no specific pattern within 100 ms of its previous spike,

which would be more likely associated with interneuronal ongoing activity. The mean firing rate of each unit was also calculated.

The above waveform features were all inspected for a potential bimodal distribution that would correspond to the two populations of cell types. Where possible, a *k*-means clustering procedure was performed with $k = 2$, in order to subclassify the units into these two putative populations, though this was done with care since the nature of the basic *k*-means algorithm is to always separate the data into the specified number of clusters regardless of the likelihood that the clusters truly exist, as discussed in chapter 3.

Once clustered into putative fast spiking and regular spiking cells based on the waveform features, the distributions of firing pattern characteristics across the two groups were compared, including the distribution of mean firing rates, the normalized population autocorrelations, and the distributions of the modal and mean autocorrelation lags. Cross correlations were performed between the spike timings of every cell pair, calculated with 0.25 ms precision over a range of ± 10 ms. 99% confidence intervals were derived from the mean results from 1,000 trials on shifted spike timings from that cell pair from a Gaussian distribution with mean zero and standard deviation 1 ms. When the observed cross correlation results surpassed the 99% confidence intervals of the 1,000 trials in either 0.5 to 4 ms or -4 to -0.5 ms bins of the cross correlation then a putative monosynaptic connection was inferred due to the temporal correlation between one unit firing and the other either being inhibited or excited beyond the average likelihood of finding that at any given time with their intrinsic firing rates (Figure 3.15). This was to be used in order to confirm the clustering into putative inhibitory and excitatory populations, though was done with caution as the likelihood of monosynaptic connections across separate electrodes at a minimum separation of 400 µm is unlikely, and common inputs or underlying population activity is more likely to cause such temporally confined action potentials.

# Chapter 3. Information latent in subacute microelectrode array recordings in humans

## 3.1. Introduction

A major goal in neuroscience is to explain how the complexities of human behaviour and disease may arise from the constituent components of the central nervous system. Animal models have provided many significant insights towards an understanding of these mechanisms, with the ability to manipulate the underlying components pharmacologically or genetically, and through the ability to perform recordings from single neurons via the patch clamp technique; these methods are unfeasible in humans, at least at present. As the scalability of the insights from these animal models up to human cortical processing at a systems level is unknown however, the ability to supplement the information found in animal models with recordings from humans *in vivo* is invaluable. To truly characterize the information processing mechanisms in the brain, simultaneous recordings from multiple neurons give significant insight. These recordings provide an understanding of how the many neurons in the cortex work together as networks of networks to translate the binary nature of action potentials into emergent, conscious experience.

### 3.1.1. *Microelectrode array recordings*

Advances in microfabrication techniques in recent years have resulted in the ability to create arrays of microelectrodes with ever finer recording tips (microelectrode arrays; MEAs), along with simultaneous advances in computational processing power allowing sampling rates from each of the many electrodes in excess of 10 kHz. Combined, these provide the ability to record reliably the high frequency electrographic signal of neurons near to each electrode tip, which, in arrays of many separate electrodes as used herein, have given unprecedented ability to record from hundreds of neurons simultaneously. Recordings from these devices need not focus on the activity of individual neurons (“single units”) as distinct entities. For example, the local field potential (LFP) recorded in the motor cortex of humans has been

used to predict upcoming limb movement (Perge et al., 2014), and seizures have been automatically detected a matter of seconds after onset via the activity from multiple units, without separating them into their separate neuronal sources, in a single human patient (Park et al., 2014). It is, however, likely that the activity of single units will provide considerable insight with regards to explaining the aforementioned gap in knowledge at the systems level of cortical processing; consider, for example, the extra level of information regarding position found in the phase precession of single units relative to an underlying theta oscillation (O'Keefe & Recce, 1993). An important first step is to characterize the stability of these recordings with regards to activity at the level of single units, so as to understand what information these recordings are able to provide over multiple hours, and multiple days. This in turn will be relevant in determining what information these recordings might be able to provide regarding activity in the peri-ictal time periods relative to baseline, and how this may translate to clinically relevant information.

The nature of these high sample rate, and many simultaneous, channel recordings results in very large datasets, as outlined in Chapter 2. As such, it is typical to manage the size of the data through experimental procedures that involve only recording from distinct, non-contiguous epochs, often days apart, and to assume that the single units recorded on one day correspond to those recorded on another. This assumption clearly depends on the stability of recordings from individual neurons at each electrode, given the possibilities of micro-movement of the array or neurons, and of cell death. Alterations in the subset of recorded neurons have been noted on a day to day basis, when using MEA recordings as an interface for a neuromotor prosthesis (Hochberg et al., 2006).

### 3.1.2. *Long term recordings from single units*

The stability of single units has been assessed relatively extensively in recordings from both macaques (Maynard et al., 1999; Suner et al., 2005; Barrese et al., 2013) and humans (Simeral et al., 2011; Perge et al., 2013; Masse et al., 2014; Pandarinath et al., 2015), though critically, these have all focused on long term stability in chronic implants over a matter of months to many years. As such, the aforementioned studies have all assessed the stability after an initial recovery period, with the earliest mentioned assessment occurring after 1 week in macaques (Maynard et al., 1999), and greater than a month post implantation in humans (46 days; Perge et al., 2014). Indeed, Barrese et al. (2013) described a period of

quiescence following array implantation in which isolatable single neurons gradually appear across the array over a period of weeks, and then disappear after months of recording. The majority of recordings assessed were in fact performed many months after implantation, including over 5 years post implantation in both macaques (2,104 days; Barrese et al., 2013) and humans (1,975 days; Masse et al., 2014). Importantly, the recordings after an initial recovery period, from a month or more post implantation, are likely to provide a different quality in signal to those recorded during the first days and weeks after implantation due to the acute response of the tissue to a small penetrative injury.

In recordings performed in patients with intractable epilepsy during the monitoring prior to resection, the time frame is limited to the first days, up to a few weeks post implantation, rather than multiple months. There is little information relating to the stability of isolatable single units during this acute/subacute time frame, as provided by their use in presurgical seizure monitoring, and so the first results described and discussed in this chapter characterize the drift that can be found in these recordings from distinct, brief epochs, hours apart. Subsequently, the drift in these recordings is further characterized during a longer, single 30 minute epoch per patient to gain insight into the variability that is plausible during baseline periods within individual units during this early time after implantation.

### 3.1.3. *Cell type specific activity*

An important consideration for interpreting these data is from what class of neurons these recordings arise. When analyzing spike timings from single units, the inability to confirm cell type during these recordings, or post hoc in the case of human studies, needs to be taken into account. The neocortex is comprised of roughly 20 percent inhibitory, GABAergic interneurons, and 80 percent excitatory, glutamatergic pyramidal cells (Douglas et al., 2004). It is reasonable to anticipate that the cell type specific responses to the same stimulus would be quite dissimilar in many situations, given their different roles in network activity, their different cellular properties (Freund & Buzsàki, 1996), and how they are “wired” within the network (Pfeffer et al., 2013). As such, drawing insight from the activity of individual neurons in these recordings without regard to their resultant effect on the network could likely be often misleading. This discrepancy is highly relevant to the work I present here, and subsequent similar recordings in pre-surgical monitoring in epilepsy patients, due to the

wealth of information that sub-classifying by cell type would provide regarding the nature of inter-ictal events, seizure onset, propagation and seizure termination.

Sub-classification of single units by cell type has been encouragingly successful in animal recordings, by utilizing underlying characteristic differences in waveform and firing properties. Importantly, neurons which are fast spiking; a property deemed typical of the somatic targeting parvalbumin GABAergic interneurons, show quicker spike kinetics due to the prompt activation of $K^+$ ion return currents (Connors & Gutnick, 1990; Skaggs et al., 1996; Martina et al., 1998; Csicsvari et al., 1999; Erisir et al., 1999). Similarly, a subset of regular spiking cells show inter spike intervals and autocorrelation patterns characteristic of neurons with a tendency to fire in successive, quick bursts – such as the $Ca^{2+}$ spike induced bursts of a pyramidal cell (McCormick & Feeser, 1990). Therefore, putative sub-classification of pyramidal cells and interneurons may be successful based on electrophysiological properties. Classifications utilizing these discrepancies have been confirmed in animal models by a variety of studies with dense enough arrays to assess cell to cell interactions: cross correlations of cellular activity indicative of supposed monosynaptic inhibition and excitation have been used to confirm these prior putative separations (Barthó et al., 2004). Further techniques, such as juxta-cellular labelling of the recorded cells have provided anatomical confirmation also (Klausberger et al., 2002).

Whilst successful separation of cell classes based on spike sorting has been used in animal studies for over a decade, such sub-classification is still rare not only in human studies, but also in recordings from the neocortex rather than from hippocampal depth electrodes. These sub-classification techniques have been touched on briefly in neocortical recordings from humans (Peyrache et al., 2012), but an in depth assessment of the plausibility of cell type sub-classification in humans, from multiple patients, has not been performed. This chapter therefore closes with an analysis of multiple metrics that may be used to attempt a putative sub-classification into fast spiking, presumptive interneurons, and regular, though burst spiking presumptive pyramidal cells, across all four patients studied here.

## 3.2. Results

### 3.2.1. *Stability of electrode specific features over 48 hours*

As outlined in section 2.2.2, the plot of the first two principal components from the collected waveforms from a given electrode has a distinctive distribution – a kind of electrophysiological signature. An example of this method can be seen in Figure 3.1 A and B, wherein the waveforms that surpassed the threshold for detection from a single electrode have been plotted from 3 minute epochs, 6 hours apart (A). For the purpose of illustration, waveforms that corresponded to a subsequently defined single unit have been shown in blue, with the background activity from more distal, non-isolatable cells (multi-unit activity, MUA), shown in dark grey. In Figure 3.1 B, the first two principal component scores have been plotted, using the same colours, highlighting the relationship between the isolatable waveforms of a single unit, and the clusters of points in principal component space. Note that the axes in principal component space are maintained throughout all epochs, with the principal components calculated as a group from all epochs at once, and so the scores from one epoch are directly relatable to those in another.

Because many such signatures were being collected simultaneously from multiple electrodes, the between-electrode variance found in these principal component space plots could be used to assess the temporal stability of the recordings, with direct relevance to the potential for isolatable clusters of single units. The between-electrode comparisons are by definition comparisons of different units as the electrode tips were separated by a minimum of 400 µm, and so provide an estimate of the variance of the electrophysiological signatures.

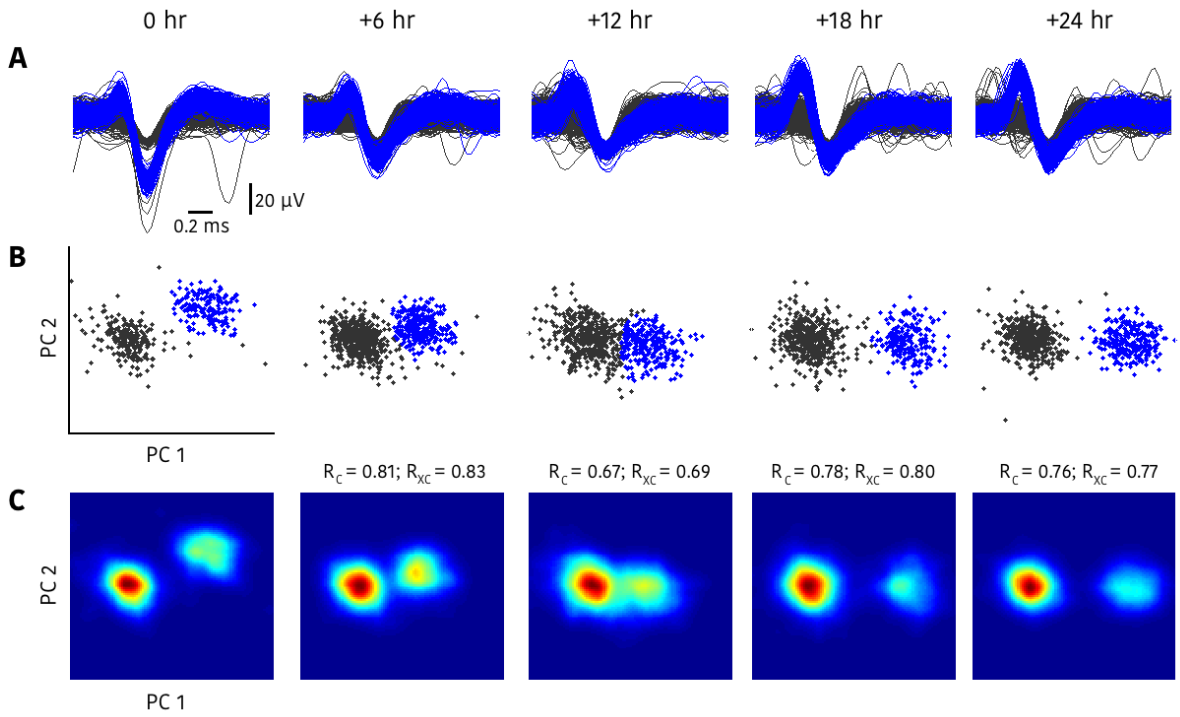


**Figure 3.1. An example analysis procedure, from spike waveforms to electrode features**
**A.** Detected waveforms from a single MEA channel from 180 second epochs, taken every 6 hours. For clarity, a subsequently defined single unit is highlighted in blue, with the background activity from more distal cells or chance deflections that reached threshold for detection in dark grey. Note the change in spike shape for the unit in blue over the 24 hour period. **B.** Corresponding first 2 principal component scores for the epochs shown in A, with the cluster of scores corresponding to the defined single unit also shown in blue. The principal component space axes are maintained throughout all epochs, with the principal component scores having been calculated across all waveforms simultaneously, and subsequently re-separated out into their separate epochs, to allow for direct comparison. Note the drift of the blue cluster relative to the background activity in dark grey, with the separation between the two corresponding to the qualitative separation of the two groups of waveforms seen in A. **C.** Smoothed probability density histograms of the electrode features seen in B, as explained in Figure 2.2. These allow for quantification of the similarity between epochs, through a conversion of the point-wise data in B, to 100 discrete bins in each dimension. The similarity between the first and subsequent epochs can be first seen through the correlation coefficient ($R_C$), which compares the density from one location with the exact same location in the comparison epoch. Further understanding of the movement of wave features can be found in the maximal cross correlation coefficient ($R_{XC}$), which is the maximum correlation coefficient found amongst every possible overlap between the two density histograms. Seen here, $R_{XC}$ is slightly larger than $R_C$ throughout, suggestive of slight drift of all principal component scores from one epoch to another, beyond the movement between the blue and dark grey clusters.

To quantify these comparisons, a metric from these distributions must be calculated. Since it is alterations in the dispersion of these points that correspond to alterations in wave shape, I first calculated the distributions of the average distance between points in principal component space to their nearest neighbour. At the same time, the overall mean distance between all points in principal component space was calculated – the pairwise distance. Neither of these methods, however, corresponded directly to the qualitative results seen in the dataset. In order to illustrate this issue, Figure 3.2 shows the resultant values from these metrics over a time period which included a seizure, during which the firing rate of single

units, and their extracellular wave forms can be expected to alter, along with the corresponding distributions in principal component space. The final panel in Figure 3.2 A corresponds to an epoch containing a seizure. In A, we see an example channel, with the nearest neighbour values shown above in blue ("N.N."), and the mean pairwise distance between all points shown in purple ("P.Dist."). As can be seen, the nearest neighbour value did not increase during the seizure epoch relative to the baseline epochs, despite the intense alteration in principal component score distribution, and whilst the mean pairwise distance did indeed show the anticipated increase, an equivalent increase was also found during non-seizure epochs, which did not show as much alteration. In B, we see the equivalent results from all assessed electrodes, with the same issues as mentioned in the two metrics arising also at the population level.

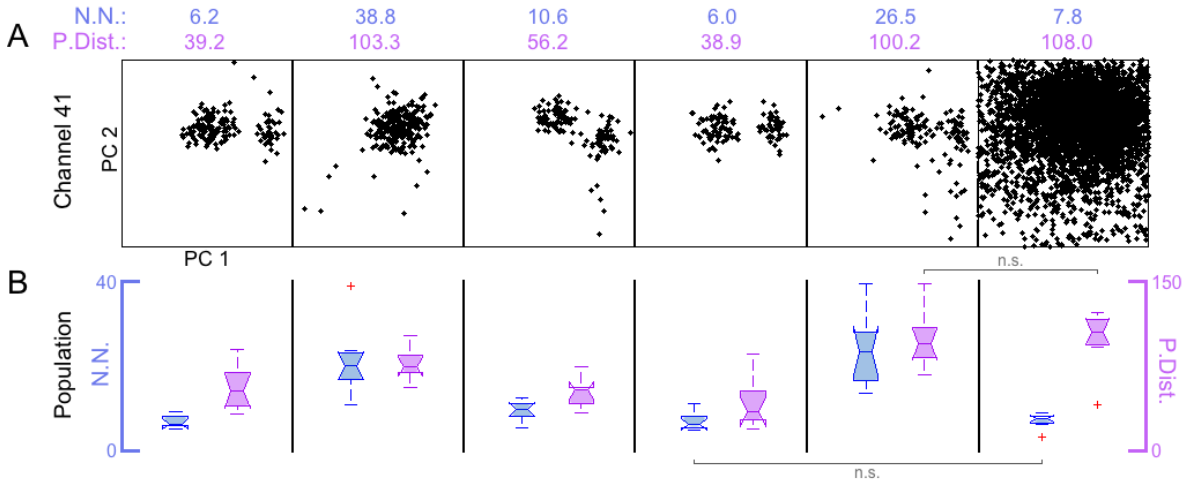


**Figure 3.2. Nearest neighbour and pairwise distance metrics**
**A.** Example results from both the mean nearest neighbour ("N.N.", blue) and mean pairwise distance ("P.Dist.", pink) methods, over multiple from a single channel, including an epoch that contained a seizure (rightmost panel). Note that in the nearest neighbour method, which is the average of the minimum distance in principal component space between each point and its closest neighbour, the loss of a distinct cluster did increase the value (2nd panel from the left), but during a seizure, when the principal component space alters considerably qualitatively, no such increase is found. In the pairwise distance method, which is the mean distance in principal component space between every possible pair of points, the seizure epoch did show an increase, though an increase of the same magnitude was found during baseline epochs also (2nd panel from the left, and 2nd panel from the right). **B.** The population results of the example shown in A, across all electrodes that showed evidence of single unit activity (*n* = 30 channels; N = 1 patient), showing that these features in the nearest neighbour (blue, left scale bar) and pairwise distance (pink, right scale bar) methods were maintained. There was no significant difference between the qualitatively altered seizure epoch and at least one baseline epoch, with one non-significant comparison highlighted for each method.

Although the nearest neighbour and mean pairwise distance metrics both had the advantage of utilizing more than just two dimensions in the principal component space, their lack of correlation to the distributions of clusters seen qualitatively in the dataset led me to develop

a second metric of the feature space. To do this, I calculated the probability density histograms of the data points for each electrode at each 180 second epoch, in 100 bins in both dimensions, and smoothed them with a moving average of 5 bins. These smoothed probability density histograms were normalized in feature space, such that density in one location corresponded directly to waves of the same shape as to those found in the same location in another epoch.

An example of this procedure can be seen in Figure 3.1 C, with each panel showing the smoothed probability density histogram from the equivalent distribution in principal component space seen in B. This method had the advantage of characterizing alterations in both unit waveform (the position of densities in the principal component space), and in the firing rate of the units during that epoch (the relative densities in each position). These distributions were calculated on 180 second epochs over a 48 hour period in one patient in order to characterize the potential alteration in this acute/subacute time frame over a 2 day period. These epochs were taken as close to one every hour as feasible with the dataset, on all channels that showed evidence of activity distinct from the homogeneous background multi-unit activity ($n$ = 30 channels).

Two dimensional correlations and cross correlations were performed on these smoothed plots. Firstly, these were performed between all electrodes within each epoch (red inter-electrode plots in Figure 3.3 A and B for correlation coefficient [$R_C$] and maximal cross correlation coefficient [$R_{XC}$] respectively). Secondly, the equivalent correlations were then performed also within electrode, relative to the first epoch on that channel (blue intra-electrode boxplots in Figure 3.3 A and B). This demonstrated a low similarity of different electrode recordings, an effect that was highly reproducible throughout the full 48 hour period, with the difference between the inter- and intra-electrode values being significant at every time point in both the correlation coefficient and maximal cross correlation coefficient measures (maximal $p < 0.001$ in both cases, Mann-Whitney U test, Holm-Bonferroni corrected). These low correlation coefficients in the inter-electrode results represent a large spread of possible electrode signatures. In contrast, the within electrode, temporal analyses showed high correlations from all epochs throughout the 48 hour period relative to the first epoch, with the majority staying above the range of the inter-electrode values at that time point, indicative of stable wave shapes and firing rates of units over a 2 day period.

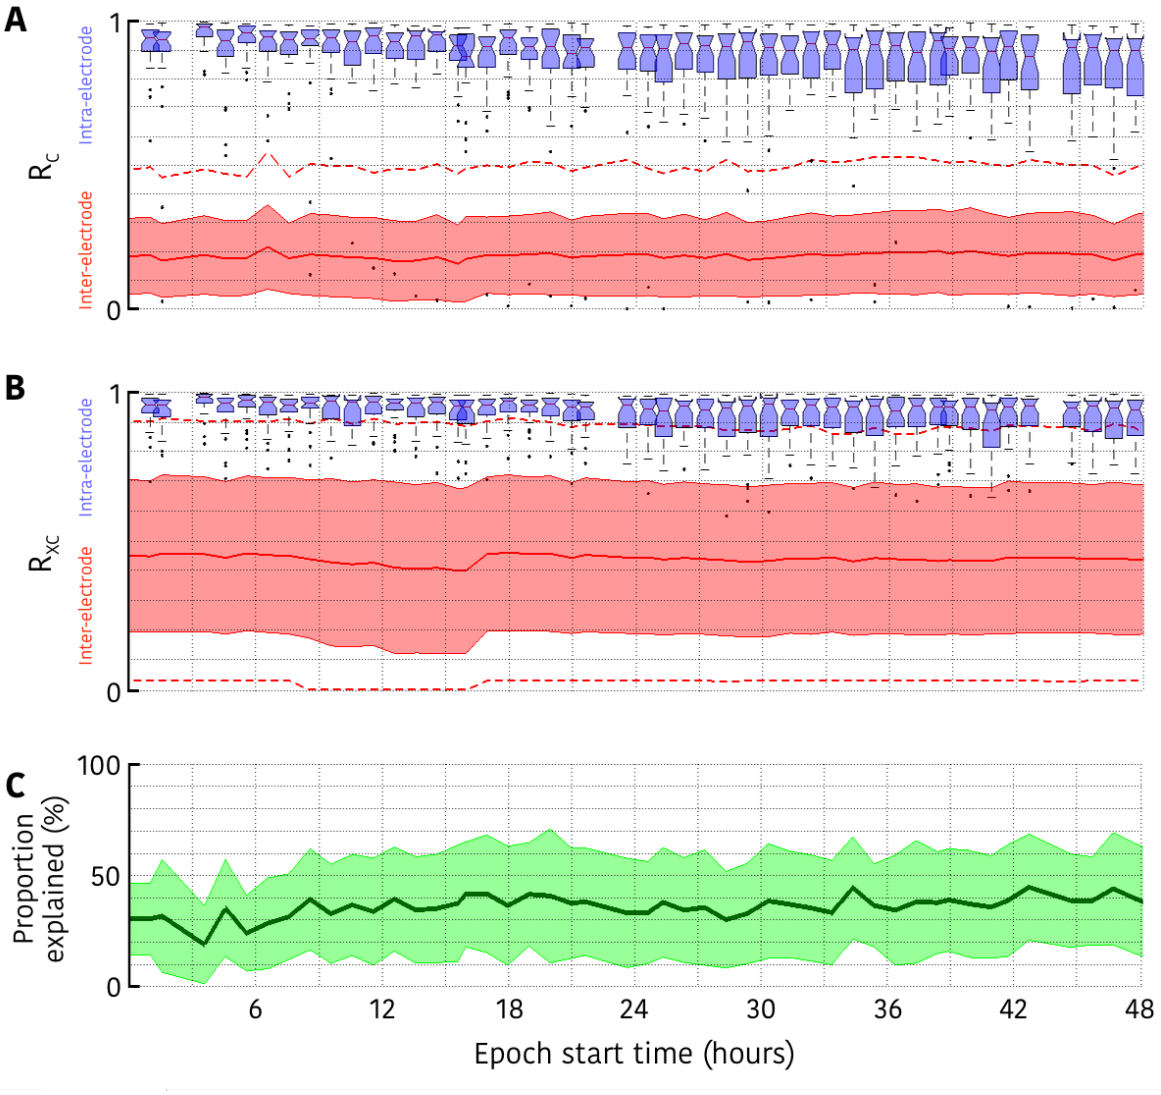


**Figure 3.3. Long term intra- and inter-electrode feature correlations**
Population results over 48 hours from every electrode that showed evidence of isolatable single unit activity during the first epoch (*n* = 30 channels; N = 1 patient), for direct correlation coefficients ($R_C$, **A**) and maximal coefficients from cross correlations ($R_{XC}$, **B**). In both, intra-electrode results are shown in the blue box plots, with the inter-electrode results in red (mean result is central red line, shading denotes ± 2 standard deviations, and red dotted lines are the range). Black dots show outliers from the intra-electrode box plots. Epochs of 180 seconds were taken as close to every hour as possible with the data available, as shown in the location of the box plots. Note the continued difference between the intra- and inter-electrode results in both methods, with the median intra-electrode result never reaching the range of inter-electrode results in $R_{XC}$ and intra-electrode interquartile range remaining higher than the inter-electrode range for $R_C$. These differences were significant throughout in both methods (maximal $p < 0.001$, Holm-Bonferroni corrected). The maximal cross correlation values through time are higher than the correlation coefficients, suggestive that while there is alteration in the electrode features, this is partly explained through a drift of the features that correspond to a single unit, rather than complete alteration of the units being recorded. The proportion of the alteration in electrode features explained by a drift of single units can therefore be measured in the difference between $R_C$ and $R_{XC}$ (i.e. $(R_{XC}-R_C)/(1-R_C)$), as shown as a percentage through time in **C.**

Both two-dimensional correlations and cross correlations were performed. In the former case, this provides a measure of how similar the two distributions are, without allowing for any leeway for movement of the densities, thereby quantifying exactly how similar the

waveforms are from one epoch to another. The maximal coefficient from the latter assessment, however, allows for a drift in the feature set, so long as its distribution relative to itself is maintained regardless of its overall location within the feature space. As such, the difference between the correlation coefficients and the maximal cross correlation coefficients provides an insight as to how much any alteration through time corresponds to a steady drift of the same single unit features, rather than a sudden complete alteration to the features recorded by that electrode. Notably, whilst the difference between the 2$^{nd}$ and 48$^{th}$ epoch's correlation coefficients (both calculated relative to the 1$^{st}$ epoch) is a significant reduction ($p$ = 0.0056), the comparison of the equivalent epochs in the maximal cross correlation coefficients is not significant ($p$ = 0.240). In fact, the difference between the 48$^{th}$ epoch's correlation coefficients and maximal cross correlation coefficients (i.e. the difference between the two methods) was significant ($p$ = 0.012). Importantly, this is indicative that there is an anticipated significant alteration in electrode features over a 48 hour period, but this alteration is predominantly a result of drift of the same distribution of features, and so utilizing recordings from epochs from between multiple days may allow for confidence in the equivalence of units over prolonged periods.

To assess how much the alteration of features through time was a result of this drift of similar features, the percentage of the drop in raw correlation coefficient that could be explained by the drift found in the maximal cross correlation coefficient was calculated through time (Figure 3.3 C). That is to say, what percentage of the loss in correlation coefficient was removed when allowing for drift in the maximal cross correlation coefficients? This showed that the level of drop in coefficient that could be explained through allowing for a drift of features corresponded to a mean of 36.2% through time, with a standard deviation of 7.4%. This percentage of the loss of correlation over time can therefore be ascribed to a drift of equivalent features, rather than a complete alteration in the principal component space.

### 3.2.2. *Single unit stability over 24 hours*

To confirm that the stability found in these electrode features corresponded to the ability to follow isolatable single units stably through time, I then performed spike sorting on a 24 hour period in patient 1, where brief (180 s) epochs were analysed every 6 hours. For these analyses, any epoch in which there was seizure activity was excluded. 76 single units were

identified in total, and of these, 54 units (71%) maintained their waveform characteristics throughout the non-seizure epochs during this 24 hour period, four units (5%) were lost and did not return, and 12 (16%) were not visible during the first analysed epoch, but appeared and persisted in subsequent epochs. There was no particular association of the appearance of new units and seizures.

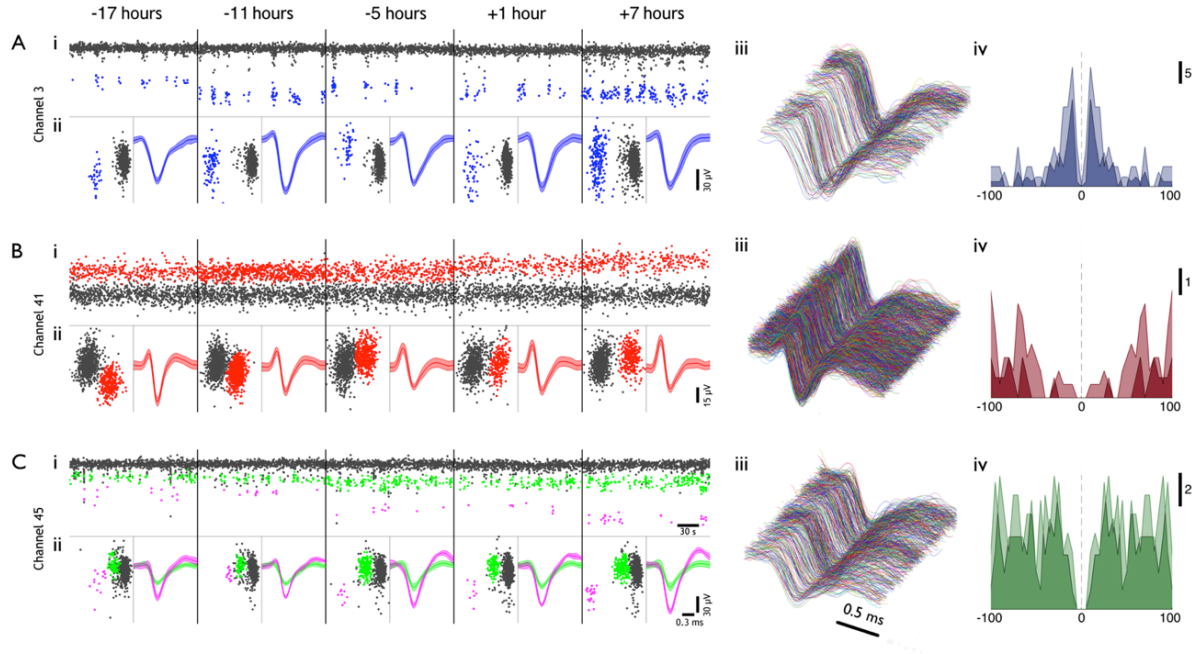


**Figure 3.4. Stability of single unit features over 24 hours**
**A – C.** Typical features from three channels in Patient 1, from 180 second epochs every 6 hours, over 24 hours, during which time there occurred a seizure. Times represent time relative to seizure onset. For each channel multiple features are shown. (**i**) The first principal component versus time; units showing clean separation from the background noise of more distal units (grey) are highlighted in colour (blue, red, green and pink). (**ii**) Left: Full epoch's first versus second principal component scores, and right: mean waveform (± 2 SD) of all spikes in highlighted cluster. (**iii**) Every waveform from each epoch plotted contiguously (only the 'green' unit in C is depicted). Note the evidence for stability of both a bursting cell, in A, and for a highly active cell, in B. While there is a drastic change in mean waveform over the 24 hours in B, the maintained distinct cluster in B(ii), without the presence of other well separated activity, implies strongly that this is a cell showing drift relative to the electrode tip. (**iv**) Autocorrelograms from time points -5 and + 7 hours relative to seizure (light and dark, respectively) from the same cells as (iii) over ± 100 ms, using 5 ms bins. Note the maintained intrinsic firing properties within each cell, and different properties between cells.

One unit was visible only in a single epoch, and five units (6.5%) were lost temporarily, for one or more epochs, but became evident again by the end of the 24 hours. Figure 3.4 shows an example 3 single units that were followed throughout the 24 hour period, showing every 6 hours for illustrative purposes. I noted instances where one unit changed wave shape while others remained stable on the same electrode (Figure 3.4 C; note the stability of the green unit while the pink unit increases in amplitude considerably). Notably, the firing patterns shown by individual units provided further confirmation of the stability of these

recordings: units displayed a range of firing features, including intermittent bursting (Figure 3.4 A), and fast spiking patterns (Figure 3.4 B), and these identifying patterns also persisted over time (Figure 3.4 A-C iv).

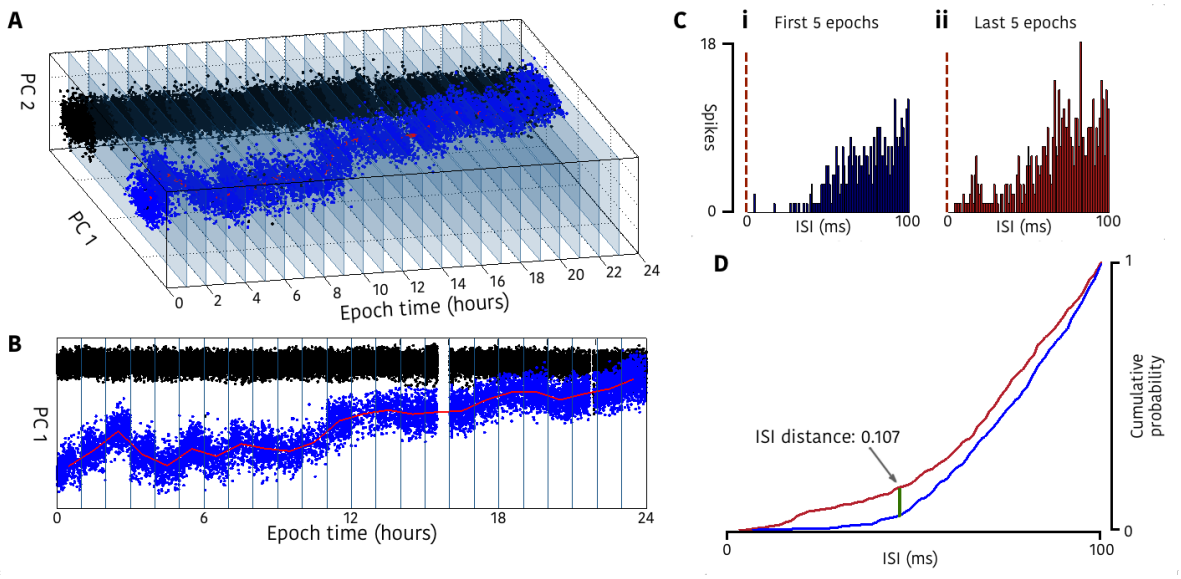


**Figure 3.5. Single unit stability metrics**
**A.** First and second principal component scores from an isolated single unit (blue), and the background multi unit activity (black) through time, with 180 second epochs taken every hour plotted contiguously. Translucent boundaries denote the edges between epochs. The centroid of the single unit cluster from each epoch shown in blue is plotted in red. **B.** Same format as for A, though with the first principal component score alone plotted through time for ease of viewing the single unit cluster centroid drift through time (red line). **C.** Histograms of the inter-spike intervals (ISI) from all spike times from the unit in blue in A & B, from the first 5 epochs (**i**) and the last 5 epochs (**ii**), over the first 100 ms. Scale is maintained across both histograms. **D.** Cumulative probabilities over the first 100 ms for the ISI histograms in C (first 5 epochs in blue, last 5 epochs in red). The difference between these cell intrinsic firing properties was quantified as the maximal difference between these cumulative probabilities as shown, referred to as the "ISI distance".

To quantify these properties at the population level, the analyses were extended to include a brief (180 s) epoch from every half hour over the 24 hour period (28 units were discerned from 20 channels) in order to follow their drifts more accurately. The drift in principal component space of the centroid from each unit's cluster was calculated through time (Figure 3.5 A & B, red line denotes centroid of cluster in that epoch), along with concomitant alterations in the mean amplitude of each single unit at each time point. The normalized drift in principal component scores over the full 24 hours was found to have a mean of 0.306 and a standard deviation of 0.231. A drift value of 0 would correspond to no movement, and a value of 1 would correspond to the centroid having moved the entire distance between its

original location and the centre of principal component space. These values were contrasted with a bootstrap estimate from the population by calculating the mean drift between a given unit and all the others in feature space, which provided a mean of 1.04 and a standard deviation of 0.653 (intra- versus inter-electrode $p << 0.001$, Mann-Whitney U test). The mean absolute change in amplitude over the 24 hour period corresponded to an alteration of 5.14 μV/hour (6.93 μV/hour SD).

The identifying nature of the firing patterns from each unit as described above was quantified to assess primarily whether the units tested at the end of the 24 hour epoch were likely the same units as those at the beginning, and also to characterize the likelihood of alterations in firing patterns from single cells. To do this, the inter spike interval (ISI) distributions from epochs that corresponded to the first 5 hours were contrasted with those from the epochs that corresponded to the final 5 hours (Figure 3.5 C). To quantify alterations in these distributions, the maximal distance between their cumulative probabilities up to a 100 ms lag was calculated (Figure 3.5 D; Suner et al., 2005), resulting in a mean ISI distance of 0.267 (0.131 SD). As per the principal component drift values, these were also contrasted to the mean inter-unit results, which resulted in a mean ISI distance of 0.360 (0.167 SD; intra- versus inter-electrode $p = 0.010$, Mann-Whitney U test).

A final measure of the stability of these units over the 24 hour period analysed was provided by the alteration of the signal to noise ratio (SNR) through time. This corresponds to the mean amplitude of the spikes from a single unit, divided by twice their standard deviation over the complete waveforms. The distributions of the SNR from all 28 units are shown through time in Figure 3.6 A. It is important to note that these calculations of SNR over time correspond only to the single units that had been followed since the beginning of the 24 hour period, and not any newly isolatable units found during that epoch, and so are not representative of the overall SNR of all single units at a given time point. Rather, instead, these values correspond to the alteration in SNR for any given unit over time, once isolated and followed continuously.

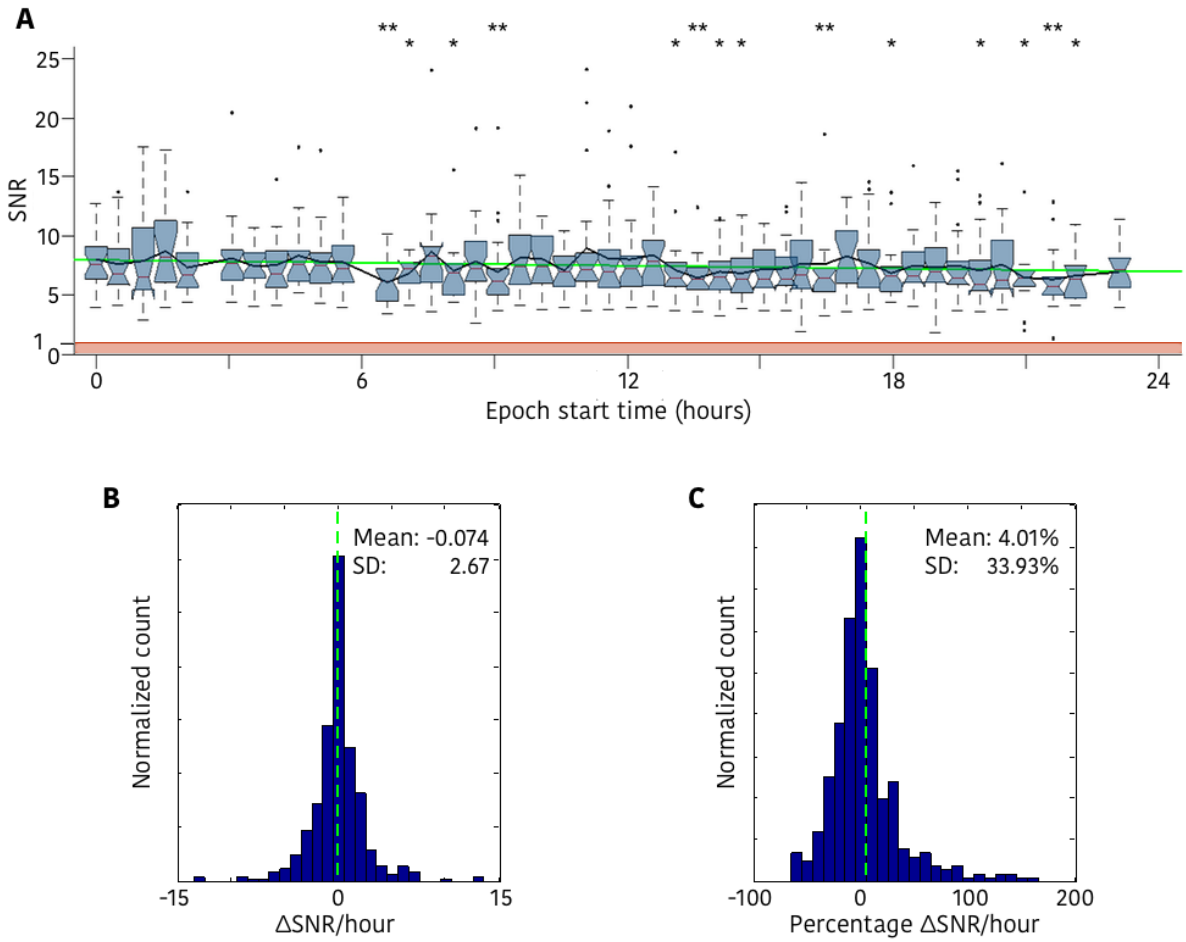


**Figure 3.6. Single unit signal to noise ratio alteration over 24 hours**
**A.** Population signal to noise ratio (SNR) from all single units that were isolated in the first epoch for this long term analysis (*n* = 28 units; N = 1 patient), from 180 second epochs taken as close to once every 30 minutes as possible in the data available; short breaks in recordings were necessitated by the patient unhooking in order to use the bathroom. The black line shows the mean of SNR through time and the green line is a linear fit to the full dataset, corresponding to a -0.04 change in SNR per hour. Black dots denote outliers from the boxplots of SNR from all units through time. The red line, and the shaded region below it, shows the level where the amplitudes of the single unit would be equivalent to the amplitude of the background noise. **B.** Distribution of all absolute alterations in SNR found for every unit, between epochs 1 hour apart, showing stability within unit, with a mean alteration of -0.074 per hour. **C.** As for B, though calculated as a percentage change in SNR rather than absolute alteration, showing very good stability on an hourly basis, with an average increase in SNR of 4.01% per hour.

Note that for the majority of the units, the SNR remains high over the full 24 hour period (Figure 3.6 A). Although there are a few exceptions where the drop in SNR was significant (likely due to low sampling from a single unit over the brief, 180 second epoch), the slope of a linear fit to these distributions would suggest only a -0.04 drop in the SNR per hour (Figure 3.6 A). A histogram of all alterations in SNR per hour showed no significant deviation from zero change (Figure 3.6 B; mean: -0.074; SD: 2.67), and when calculated as a percentage change per hour, these showed a positive mean change of 4.01% per hour (33.93% SD;

Figure 3.6 C). These showed that the SNR of single units through time over a 24 hour period was very stable in this subacute time period.

### 3.2.3. *Isolatable single units from continuous 30 minute epochs*

As the aforementioned analyses were all performed on brief epochs of only 180 seconds, I next asked if there might be fluctuations in this stability during a continuous 30 min epoch, in multiple patients, and as a result, what the quality of single units isolated during this subacute time frame may be. To do this, I performed spike sorting analyses on a continuous 30 minute epoch for each patient, and calculated in depth quality metrics on the resultant single units, as per Hill et al. (2011). In doing so, I was able to isolate a total of 361 putative single units successfully from a subset of channels in all patients (patients 1 to 4: 111; 110; 61; and 79 single units respectively). Three example units can be seen in Figure 3.7 A – C.

| Patient | Mean SNR | ± SD | Minimum | Maximum |
|---|---|---|---|---|
| Patient 1 | 6.26 | ± 2.18 | 3.56 | 17.42 |
| Patient 2 | 7.87 | ± 2.19 | 3.57 | 17.04 |
| Patient 3 | 7.66 | ± 1.94 | 4.02 | 14.96 |
| Patient 4 | 6.37 | ± 2.12 | 3.41 | 15.05 |

**Table 3.1. Signal to noise ratio (SNR) of single units in a continuous 30 minute epoch across patients**

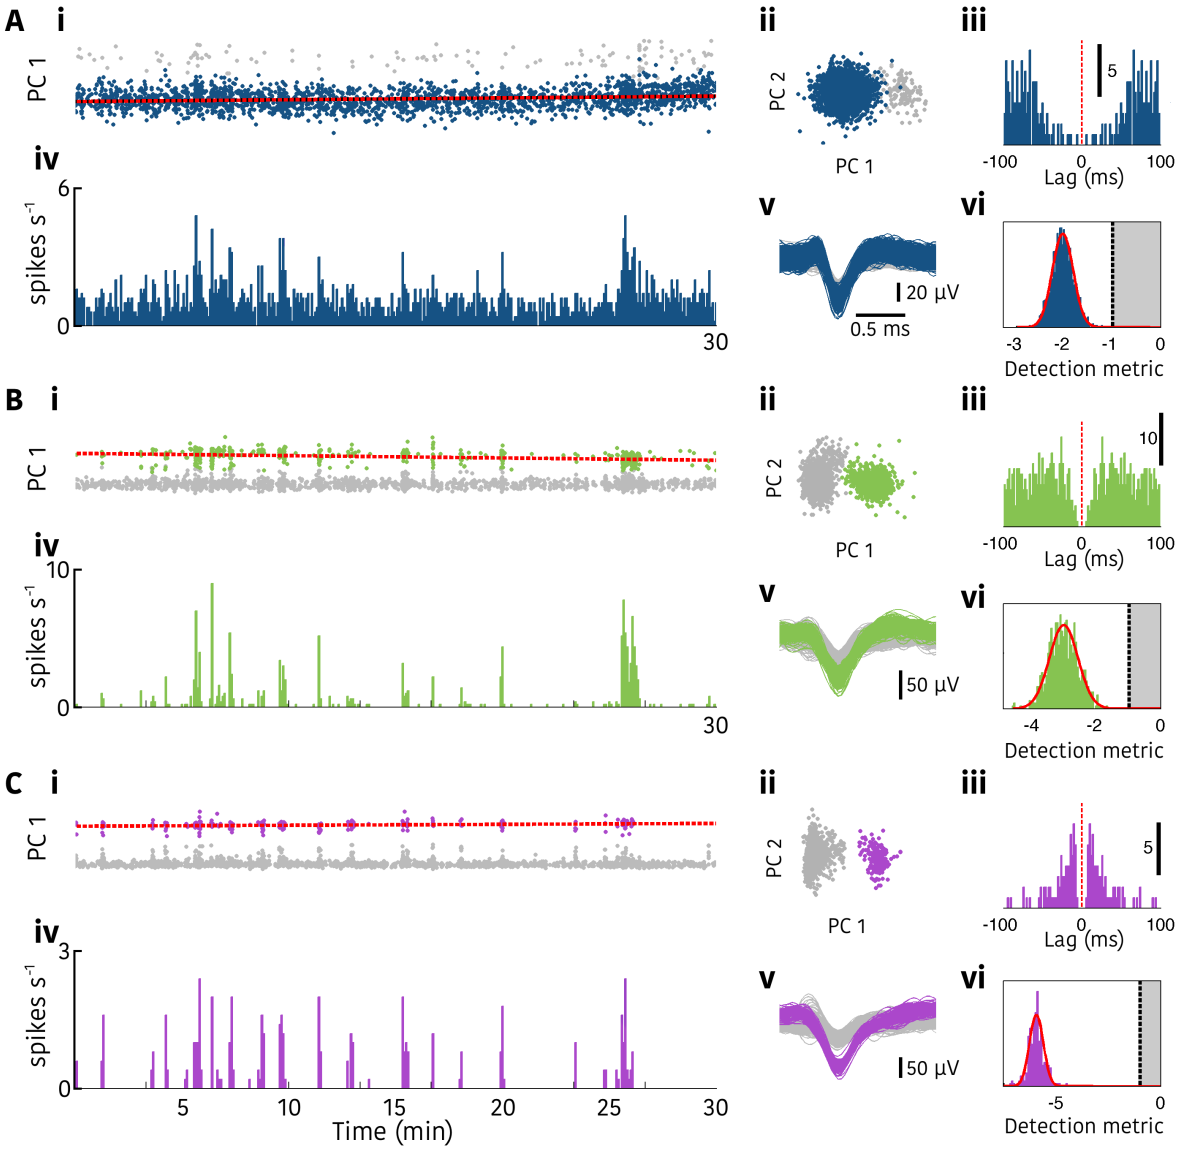


**Figure 3.7. Example single units and associated metrics from a continuous 30 minute epoch**

**A – C**. Three example single units isolated through spike sorting of a continuous 30 minute epoch, and their associated metrics, with the same format for each. The single units are highlighted in blue, green, and purple respectively. (**i**) First principal component scores through time (single unit in colour, un-clustered activity in grey). Red dashed line denotes a linear fit to the single unit scores. (**ii**) First versus second principal component scores from the full 30 minute epoch. (**iii**) Autocorrelation of all spike times from the single unit, over ± 100 ms, calculated in 2 ms bins. Note the absence of activity in the ± 2ms range, and the distinct activity patterns found in each unit shown. (**iv**) Firing rate over the 30 minute epoch for each unit, calculated in bins every 5 seconds. (**v**) All waveforms from the single unit (coloured), with background, un-clustered waveforms in grey, over a 1.6 ms window. (**vi**) Histograms of the voltage at the detected trough of each waveform in the cluster, divided by the threshold for detection, with a Gaussian curve fitted (red). The voltages at detection follow a normal distribution, and so the number of spikes that were missed from a given cluster due to being sub-threshold, can be estimated by the percentage of the Gaussian fit that sits above the threshold (> -1, grey shaded region). None of these units had expected false negatives, with none of the Gaussian curve reaching the threshold level.

The SNR of these units was high enough to be confident in their corresponding to genuine individual neuronal sources, with a global mean SNR of 6.65 (2.05 SD; range: 3.41 to 17.42). This was found to be very similar for all patients (Table 3.1). The distribution of SNR values from 30 minute epochs from each patient prior to each seizure recorded is shown in Figure 3.8.

Linear fits were applied to the amplitudes of every spike for each single unit through time, which resulted in a global mean alteration in amplitude of -0.014 µV $min^{-1}$ (0.223 µV $min^{-1}$ SD). These values were similarly small in all patients (patients 1 to 4, mean ± SD µV $min^{-1}$: 0.003 ± 0.165; -0.027 ± 0.310; 0.008 ± 0.136; and -0.051 ± 0.249 respectively). Importantly, however, the $R^2$ value for these linear fits was very low throughout (maximal $R^2$: 0.049, patient 2), indicating that while the mean alteration in amplitude per minute is relatively small, there is a much larger variance in the detected amplitudes than there is an average alteration through time, over a 30 minute period.

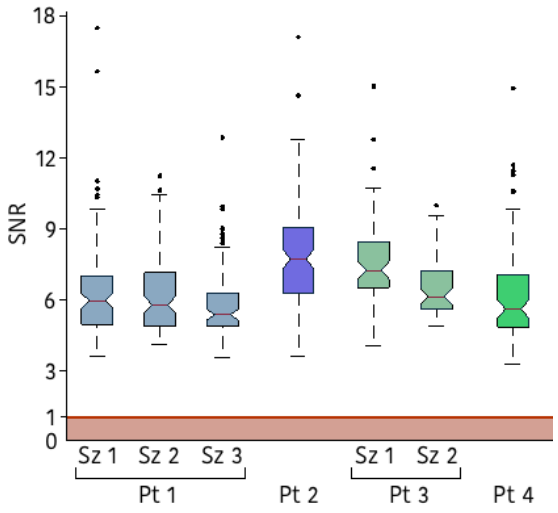


**Figure 3.8. Single unit signal to noise ratios over multiple patients and seizures**
The distribution of signal to noise ratios (SNR) for multiple separately spike sorted epochs, each of 30 minute duration immediately prior to a seizure. This included multiple seizures in patients 1 and 3, with patients grouped by colour (*n* = 111, 101, 91, 110, 61, 61 and 79 units for each seizure shown respectively). As in Figure 3.6, the red line, and shaded region below, show the area wherein the amplitudes of the single unit spikes are equivalent to, or below, the amplitude of the background noise. In reality, these spikes would not be detected and so the distribution could not fall below this line. Across all units in all patients, the minimal SNR did not drop below 3, indicative of cleanly isolated units from the background noise.

For each single unit, metrics were calculated as described in section 2.2.5, to quantify how readily accessible single unit recordings are during the acute/subacute time frame of these MEA recordings in humans. The probabilities of false negatives as estimated by the proportion of the Gaussian fit that fell below the detection threshold for their channel were very low. As described in section 2.2.5, the detected voltage for each spike varies in amplitude to result in a normal distribution due to underlying fluctuations in the signal from Gaussian noise. Gaussian curves can therefore be fitted to these detections in order to measure the expected proportion of spikes that arose from a given unit, which did not reach threshold, thereby being missed during the detection phase (single unit examples are shown in Figure 3.7 A – C vi). The resultant percentage values from this estimate of the false negatives can be seen in Table 3.2. Note the low percentages throughout, including in the maximum probabilities in each patient. Table 3.2 also includes the mean and standard deviation multiples of the detection threshold for each unit for each patient, which acts as a quantification of the separation between the unit amplitude and the background noise.

| | | Percentage sub-threshold | | | Multiple of threshold | |
|---|---|---|---|---|---|---|
| **Patient** | **Total units** | **Mean (%)** | **± SD (%)** | **Max (%)** | **Mean** | **± SD** |
| Patient 1 | 111 | 0.00038 | 0.0014 | 0.0086 | 3.23 | 1.51 |
| Patient 2 | 110 | 0.000072 | 0.00050 | 0.0048 | 4.40 | 1.61 |
| Patient 3 | 61 | 0.0054 | 0.032 | 0.24 | 2.93 | 0.99 |
| Patient 4 | 79 | 0.0024 | 0.013 | 0.10 | 2.98 | 1.48 |

**Table 3.2. Estimation of false negatives through spike detection across all units in each patient**

An estimation of the false positives per single unit cluster, as assessed by the number of absolute refractory period violations (the presence of a spike within 2 ms of a previous spike from that unit) had a global proportion of 0.00094% ± 0.003% of all spikes within the defined cluster boundaries. This corresponded to 0.0013% ± 0.004%; 0.00088% ± 0.003%; 0.00062% ± 0.001%; and 0.0017% ± 0.005% for patients 1 to 4 respectively. Across all patients, the maximal proportion of spikes within a cluster that violated the 2 ms refractory period was 0.032%. One criterion for accepting a cluster being a cleanly isolated single unit was to have less than 1% contamination of the refractory period, and thus values greater than or equal to 1% would never be expected. However, the fact that these values were considerably lower than that threshold of 1% for being discarded was indicative of single units with very

little probability of multi-unit contamination. It was also found that 94 (85% of all single units isolated in that patient), 78 (71%), 49 (80%), and 67 (85%) single units had zero spikes occurring within the 2 ms refractory period for patients 1 to 4 respectively.

It is important to note, however, that stability in one metric does not require stability in the others. There were instances in which, during a continuous spike sorting on a single epoch, sudden shifts in the waveform of a unit along with the associated principal component scores occurred, while at the same time the firing rate and ISI distribution of the unit were maintained (Figure 3.9). Note the sudden alteration in the first principal component score, roughly 60 seconds prior to the end of the epoch shown in Figure 3.9 D (blue to red), and how well separated the blue and red clusters are in Figure 3.9 A. Despite these changes, the firing rate of the blue and red units are equivalent either side of the sudden shift in amplitude, with neither cluster firing substantially while the other is present (Figure 3.9 E). At the same time, the autocorrelation histograms from the two units are very similar (Figure 3.9 C). This shift did not occur with any noticeable concomitant shift in other electrodes.

### 3.2.4. *Putative sub-classification of single units by cell type*

To assess the ability to sub-classify single units from these recordings, metrics from the cleanly isolated single units of the aforementioned 30 minute epochs from each patient were calculated, as described in section 2.3.4. For each patient, Kolmogorov-Smirnov tests for normality were applied to the distributions of the valley to peak delay of the average wideband waveform for the unit, its FWHM, the mean autocorrelation lag, and the difference between the early and late portions of the first derivative of the mean wideband waveform.

This was performed to highlight potential features on which a sub-classification into two distinct groups might be possible, because if there were a clean separation between two populations in one of these metrics, it would be expected to produce non-Gaussian distribution. As can be seen from the plots of the cumulative probabilities and empirical cumulative distribution functions (Figure 3.10), the majority of these distributions were in fact normal, with only 2 of the 16 overall metrics passing below a *p*-value of 0.05 (the derivative differences in patient 1, and the valley to peak delay of the mean waveform in patient 3).

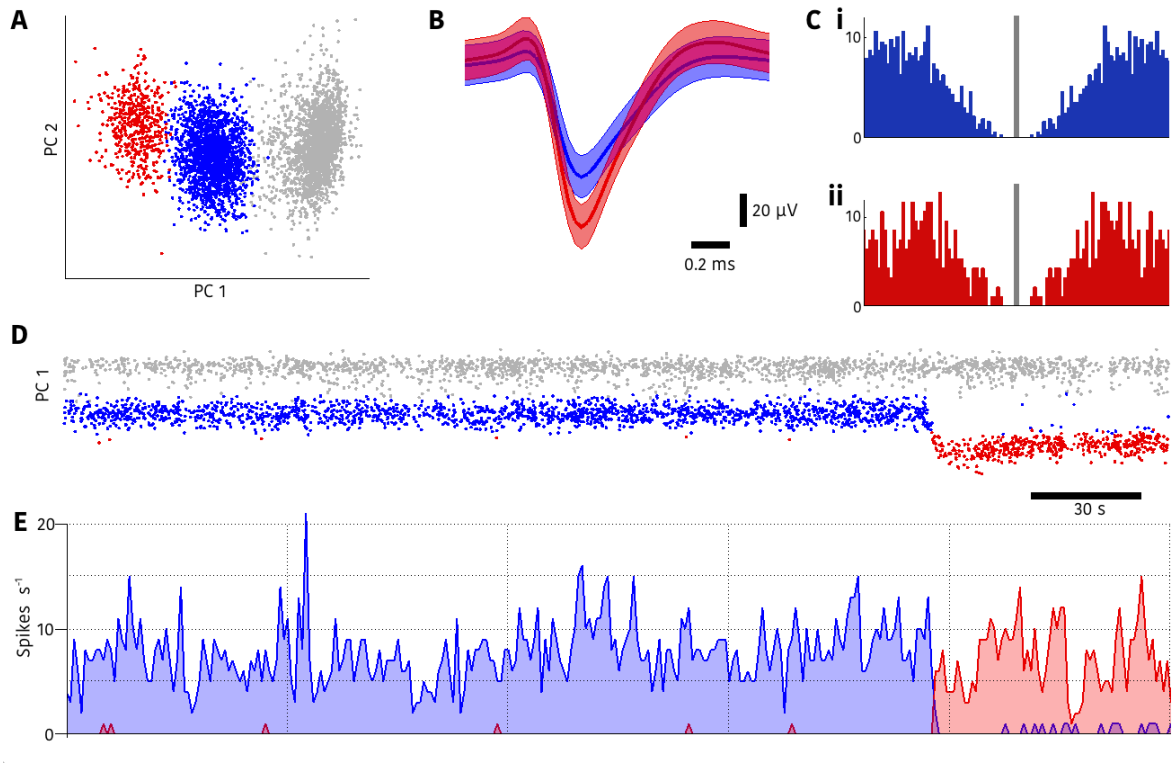


**Figure 3.9. Sudden alteration in isolated single unit waveform**

Different waveforms and the associated discrepancies in principal component scores are not necessarily evidence of distinct neurons, as shown over a single 180 second epoch from a single electrode here. **A.** First versus second principal component scores from the full epoch, showing background activity in grey, and two distinct clusters in blue and red. **B.** Mean ± 2 SD for all waveforms in the clusters shown in A, showing differences in waveforms between the two groups. **C.** Autocorrelation histograms for the blue (**i**), and red (**ii**) units. Note the similarity of this cell-intrinsic firing property between both units. **D.** First principal component scores through time over the full 180 second epoch. Note the sudden transition from one cluster to the other roughly 60 seconds prior to the end of the epoch, with very few spikes assigned to the other unit during each period. It is very unlikely that one unit should stop firing almost entirely at the same time as another starts firing with similar firing patterns (as seen in C), having been quiescent throughout the epoch beforehand. The two clusters arising from a single unit is further substantiated by the firing rate through time (**E**), which shows a very similar firing rate for both units, and near total quiescence while the other unit is active.

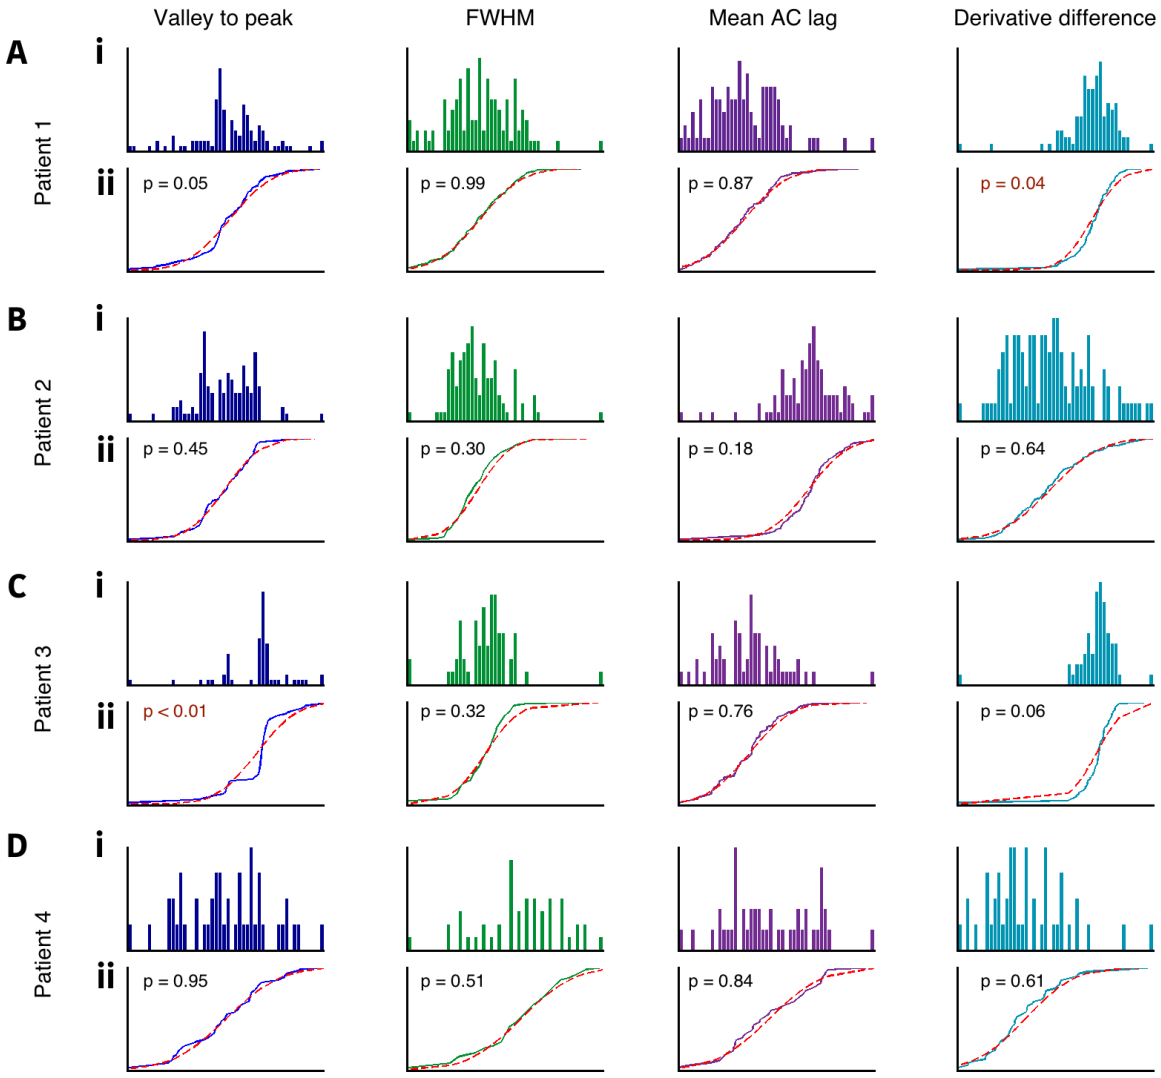


**Figure 3.10. Single unit metrics associated with distinct cell types**

**A – D.** Population features from the mean broadband waveform and firing properties from all isolated single units in each patient, including the delay from the spike minimum to its following maximum ("Valley to peak", dark blue), the spike full width at half its maximum height ("FWHM", green), the mean lag of the unit's autocorrelation ("Mean AC lag", purple), and the difference in area between the early and late components of the first derivative, as explained in Figure 2.14 ("Derivative difference", light blue). (**i**) Histograms for each metric in each patient, normalized to their minimal and maximal values on the x scale. If these metrics were able to sub-classify the population of single units into two distinct groups, they would need to form bimodal distributions. Therefore, the y scale is not of importance here, but rather the distribution. (**ii**) A bimodal distribution is necessarily non-normal, and so feasibility for sub-classifying on one metric alone can be measured by its deviation from a Gaussian distribution. Coloured lines show the empirical cumulative distribution functions for these metrics, and the red dashed lines show the expected cumulative distribution function from a Gaussian fit to these data. Discrepancies between the two highlight possible metrics on which a sub-classification might be performed. The majority of these distributions were not significantly different from a normal distribution, with the exception of the derivative difference in patient 1 (A), and the valley to peak delay in patient 3 (C), though the valley to peak delay in patient 1 was near threshold for significance (0.05), as was the derivative difference in patient 3 (0.06).

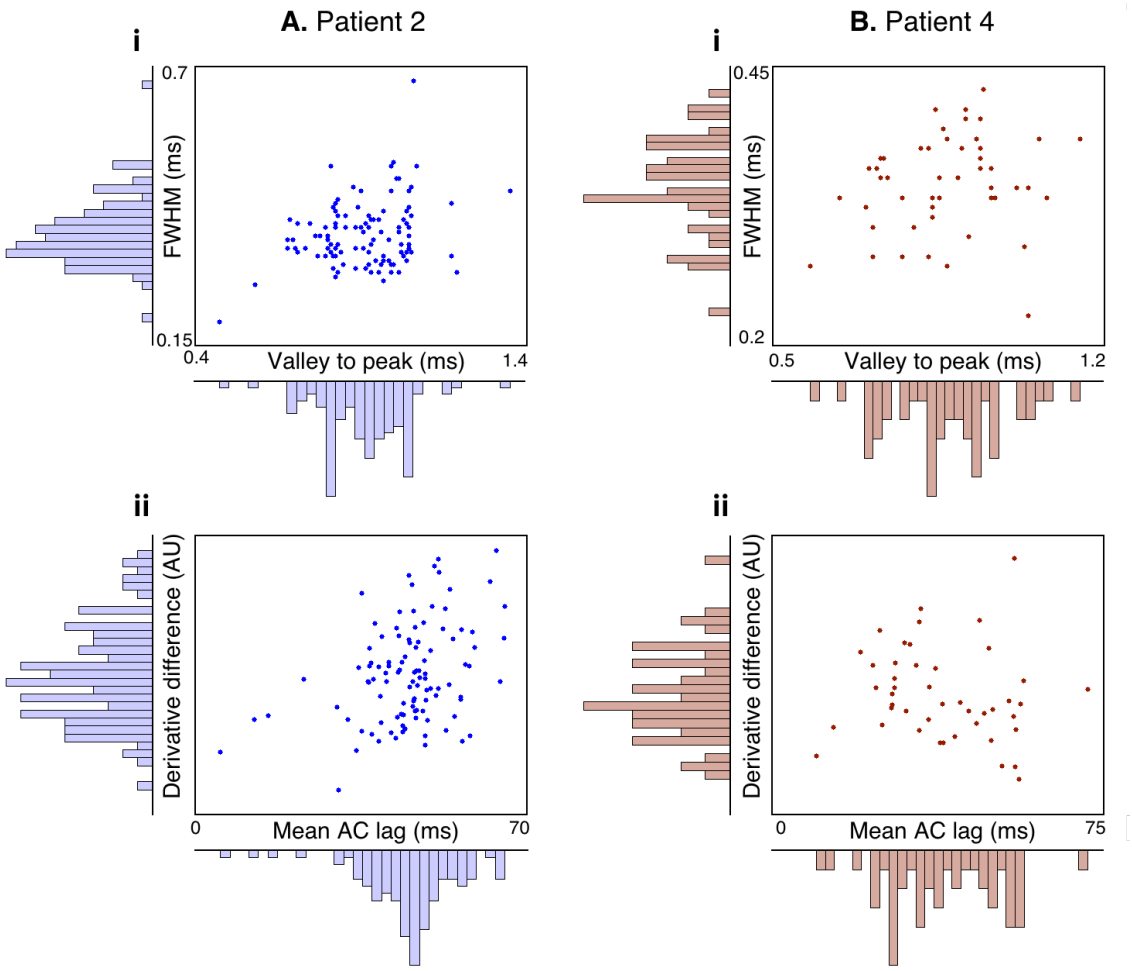


**Figure 3.11. Inability to sub-classify by cell type in patients 2 and 4**
Plots of the delay from the mean broadband waveform minimum to its following maximum ("Valley to peak") versus the spike full width at half its maximum height ("FWHM") (**i**), and the mean lag of the unit's autocorrelation ("Mean AC lag") versus the difference in area between the early and late components of the spike kinetics ("Derivative difference") (**ii**). These are shown from all isolated single units in patients 2 (**A**), and 4 (**B**), showing the inability to cleanly define boundaries between clusters arising from distinct cell types.

Whilst this was indicative that the majority of measurements were not likely to provide information on which the single units could be confidently separated into two populations, it was plausible that a combination of metrics would highlight clusters in multiple dimensions of near Gaussian distributions. In Figure 3.11 it can be seen that no clear clusters were visible in the projections of these features for patients 2 (A) and 4 (B), and so no attempts to sub-classify these single units were performed, as a *k*-means clustering with *k* equal to 2 would split these populations with no clear boundary in the feature space. The outcome of a sub-classification of the valley to peak delay and FWHM features for patients 1 and 3, along with the corresponding properties of the two populations, can be seen in Figures 3.12 and 3.13 respectively.

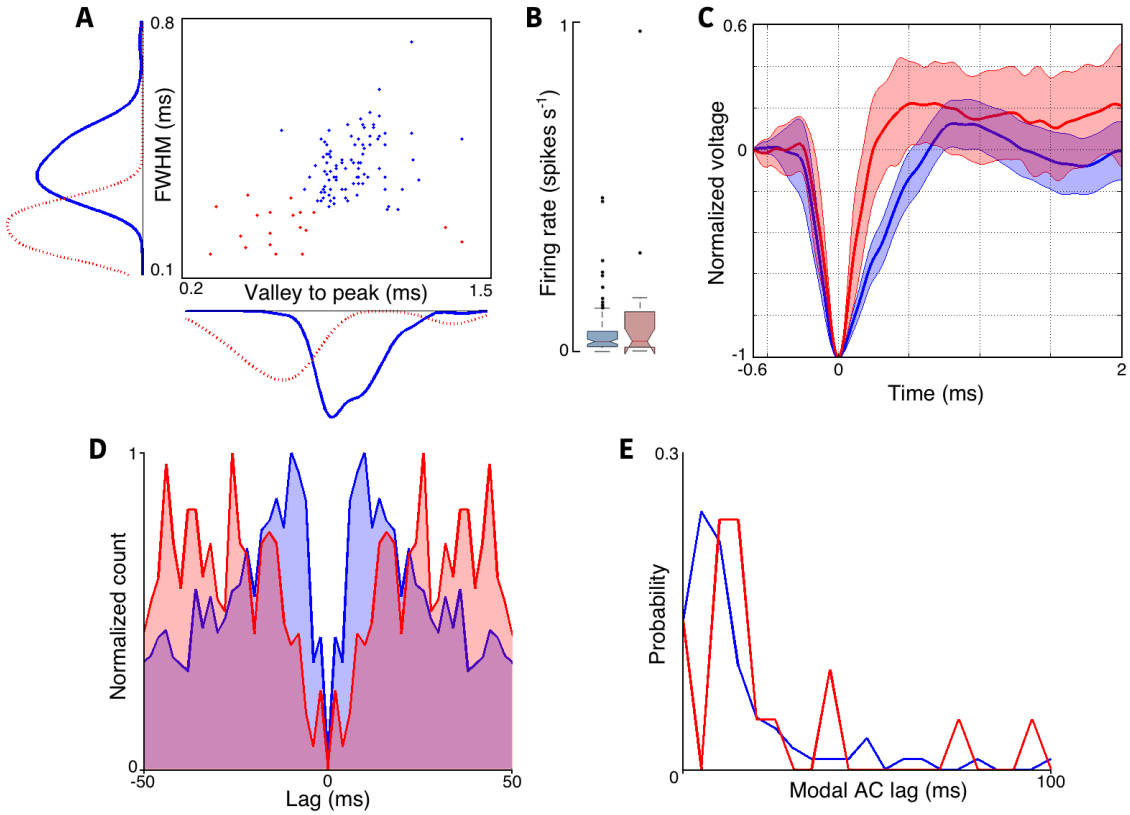


**Figure 3.12. Putative sub-classification by cell type in patient 1**
**A.** Results of *k*-means clustering on mean features from each single unit in patient 1, using the delay from the mean broadband waveform minimum to its following maximum ("Valley to peak") and its spike full width at half the maximum height ("FWHM"), with resultant probabilities in each dimension shown. Red shows the fast spiking units that are putative interneurons, and blue shows the regular spiking units that are putative pyramidal cells, with colours maintained throughout. **B.** Distribution of mean unit firing rates for the two sub populations over the 30 minute epoch. **C.** Mean ± 2 SD of all unit mean broadband (0.1 Hz to 5 kHz) waveforms for the two clusters, showing the briefer action potential of the putative interneurons in red. Waveforms have been normalized to a point 0.6 ms prior to the spike, and their minima. **D.** Autocorrelograms from the mean of all autocorrelation histograms across the two populations, normalized to their maxima, over the ± 50 ms range. These show the firing patterns associated with interneurons in red, and the greater likelihood of bursts of action potentials in the putative pyramidal cell population (blue). **E.** The probability distribution in the modal lag in the autocorrelation histograms over 100 ms for the two populations, showing a slight increase in the putative fast spiking population.

The valley to peak delay and FWHM features were used alone for this classification, as primarily, they displayed reasonably discrepant clusters for these patients, but also so as to be able to assess the accuracy of the sub-classification through the firing patterns of the populations. This subsequent assessment was done using the autocorrelation patterns and the overall firing rate of the units. The firing rate might be expected to be related inherently to the autocorrelation patterns, and so clustering also on the autocorrelation features would potentially bias the resultant firing rates of the two populations and thus the mean autocorrelation lag data were excluded for the automatic clustering.

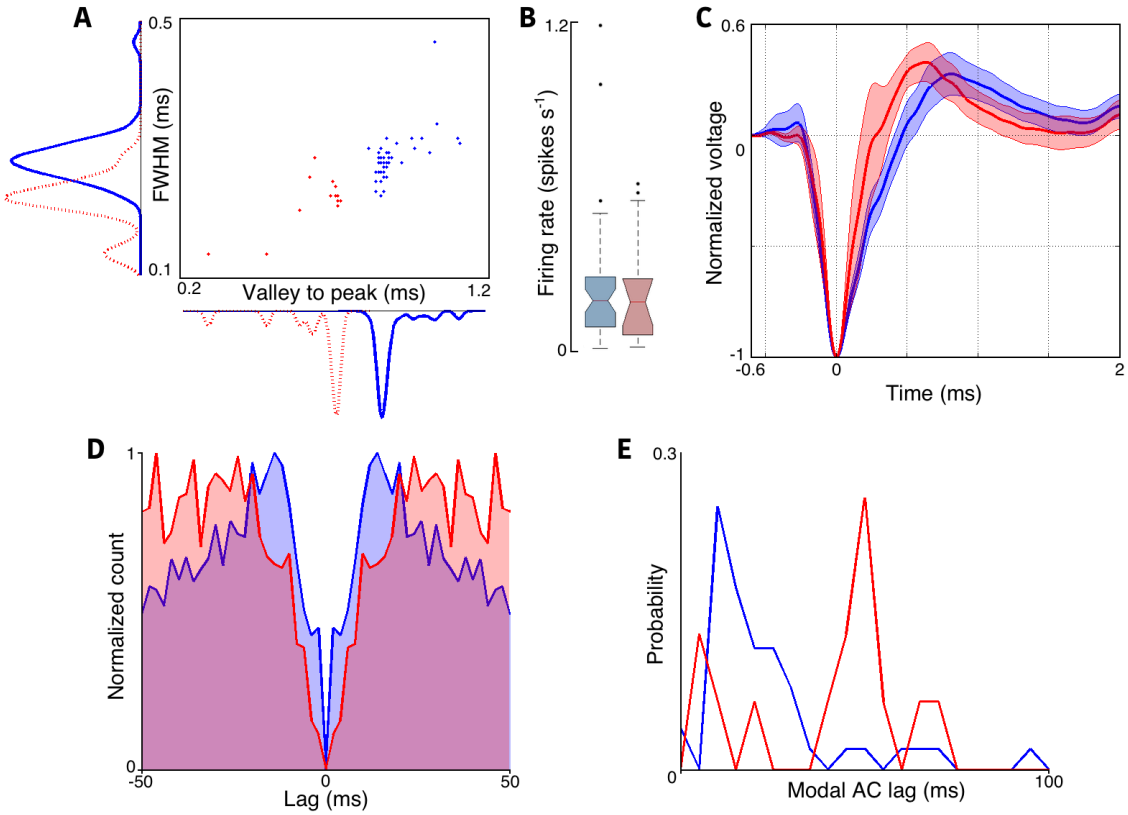


**Figure 3.13. Putative sub-classification by cell type in patient 3**
The same format as for Figure 3.12 throughout, though performed on the single units from patient 3. Note however, the lack of difference in firing rate here (**B**), though maintained alterations between populations in autocorrelation features (**D** & **E**).

The firing rate in the resultant populations from patient 1 showed the expected higher rate in the putative fast spiking interneurons, though this was not significantly higher than the regular spiking population (Figure 3.12 B). The normalized autocorrelation histograms for the two populations over ± 50 ms showed properties that were indicative of their putative cell types (Figure 3.12 D), as did the modal autocorrelation lags for the two populations (Figure 3.12 E). This sub-classification corresponded to a 17% proportion of putative fast spiking interneurons for patient 1. No expected difference in firing rate was found in the two populations from patient 3 (Figure 3.13 B), though both the normalized autocorrelation histograms and modal autocorrelation lags of the two populations were indicative of their putative cell types (Figure 3.13 D & E). This sub-classification corresponded to a proportion of 21% for putative fast spiking interneurons for patient 3.

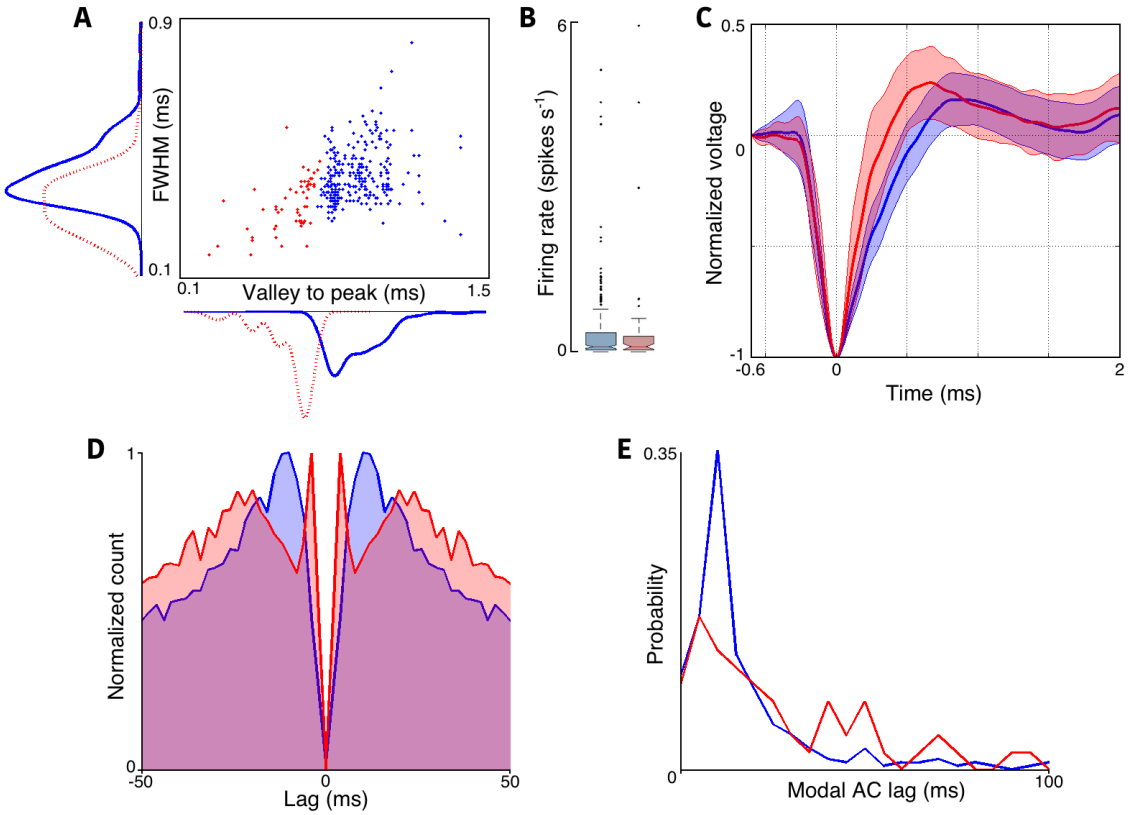


**Figure 3.14. Putative sub-classification by cell type across all patients**
The same format as for Figures 3.12 and 3.13 throughout, though performed on all single units from all patients simultaneously. Note, however, that the anticipated difference in firing rate between the two populations is not present (**B**), the difference in broadband mean spike duration is qualitatively less defined (**C**), and the anticipated differences in the cell type specific autocorrelation features are not convincingly present (**D** & **E**).

I asked whether this ability to sub-classify might be impaired by relatively low numbers of single units, as typically, such sub-classification has been performed on many hundreds to thousands of units at once in animal studies (Barthó et al., 2004). To test this hypothesis I pooled all single units from all patients, and attempted to sub-classify using the same methods on the total population (Figure 3.14), as has been shown successful in neocortical recordings in macaques (Mitchell et al., 2007). It was found that separating the features based on their valley to peak delay and FWHM, as before, produced clusters of waveforms of two distinct durations as expected. However, the firing rate was not indicative of the two putative cell types, and although there was a slight alteration in the normalized autocorrelation histogram, this was not convincing enough to reliably suggest that these were acceptably sub-classified exclusively by cell type in order to draw any further conclusions. Sub-classifying across all epochs and patients resulted in uneven, though plausible, representation across groups, with 18.0%, 10.0%, 23.0% and 22.9% of units being assigned as putative fast spiking interneurons for patients 1 to 4 respectively.

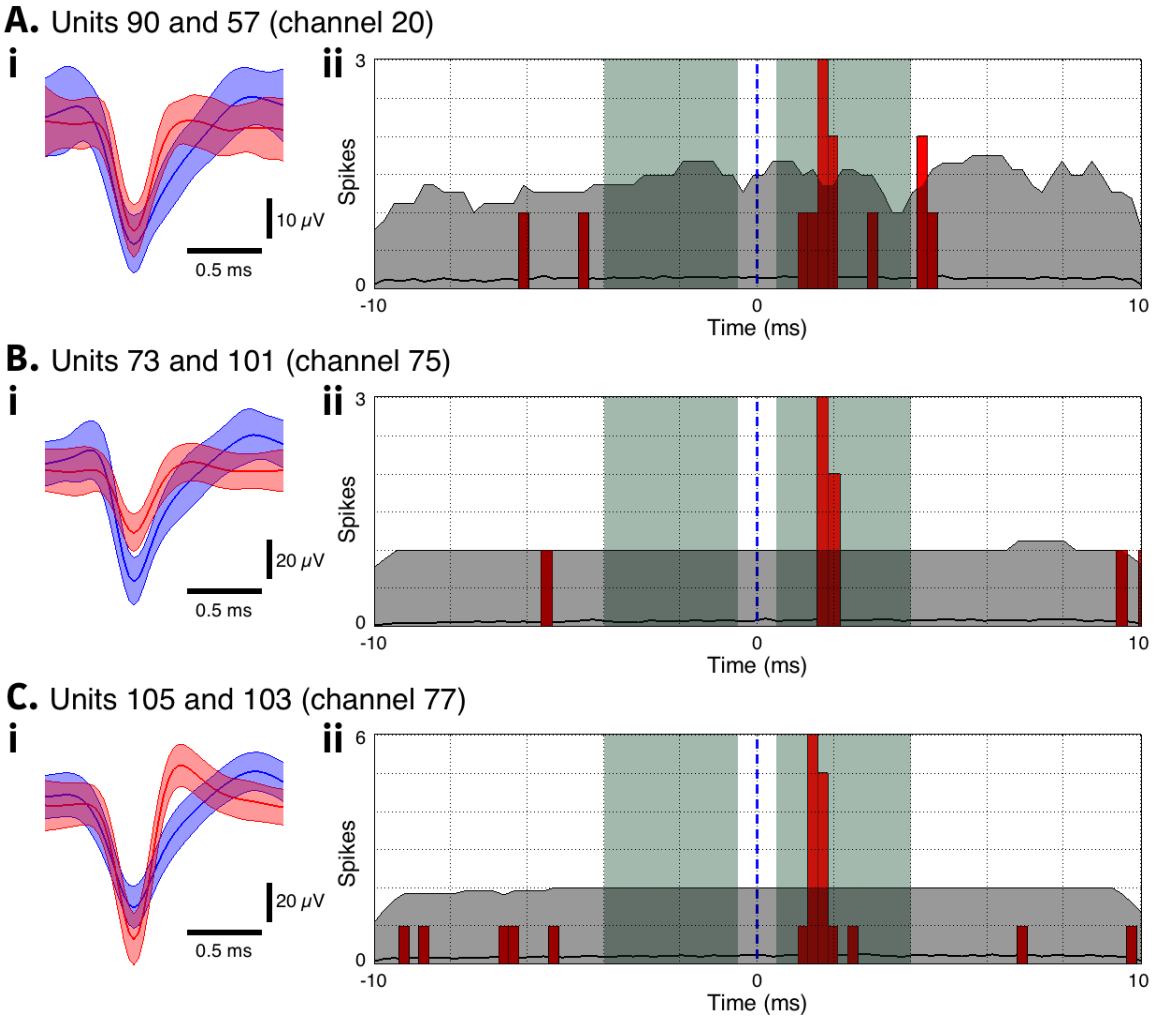


**Figure 3.15. Putative monosynaptic connections between single units**
**A – C.** Features from three putative monosynaptic connections between cell pairs on the same channel as one another that were detected in cross correlation patterns. (**i**) Mean ± 2 SD from all waveforms within the putative presynaptic (blue) and postsynaptic (red) units, showing that these arise from different neuronal sources, rather than the bursting nature of a single unit that has been incorrectly split into two clusters. (**ii**) Cross correlations, showing how many spikes occurred in the postsynaptic unit relative to when spikes occurred in the presynaptic unit, in 0.25 ms bins over ± 10 ms (red bars). 99% confidence intervals were calculated from 1,000 trials in which the spike times were each shuffled an arbitrary amount selected from a normal distribution with a standard deviation of 10 ms each time (grey shaded region, black line within the shaded region denotes the mean result of the trials). Putative connections were detected where the real cross correlation results were either above (excitatory) or below (inhibitory) this 99% confidence range for two consecutive bins, during a window of between 1.5 and 4 ms before or after a spike in the presynaptic unit (dark green shaded regions). Due to the low firing rate of units, the lower bound of the confidence intervals is zero, and so no putative inhibitory connections could be found due to a floor effect.

Cross correlations were performed between the spike times of each unit with every other, so as to locate potential significant correlations indicative of monosynaptic connections, as described in section 2.3.4. This was performed in order to confirm the aforementioned sub-classifications. In doing so, three putative monosynaptic connections were found amongst

the single units from patient 1 (Figure 3.15), though none were found amongst the single units from patient 3. All three putative monosynaptic connections were excitatory, which did not oppose their prior sub-classification by wave shape, and each was paired with another single unit from the same channel on the MEA. The percentage of spikes from the putative presynaptic cell that had a corresponding spike within 1.5 to 4 ms in the second unit was 5.20%, 7.69% and 2.40% for the pairs on channels 20 (Figure 3.15 A), 75 (Figure 3.15 B) and 77 (Figure 3.15 C) respectively.

## 3.3. Discussion

These analyses confirm previous studies of MEA recordings in humans (Hochberg et al., 2006; Truccolo et al., 2011) and animals (Suner et al., 2005; Gilja et al., 2006; Linderman et al., 2006) that these devices can isolate single unit neuronal firing patterns. These latter studies, in which the long-term stability of the units was explicitly addressed, were of chronic implants, but seizure monitoring is necessarily conducted over a far shorter time frame, over a period of days to weeks. These recordings demonstrate that persistently stable units can be found within this earlier time frame too, in humans, and thus provides the potential to probe ictogenic networks at the single cell level over multiple seizures. Furthermore, chronic implantations of MEAs in humans will typically be restricted to the primary motor cortices due to the nature of their use as neuromotor prostheses in patients with movement disabilities such as quadriplegia (Simeral et al., 2011), or amyotrophic lateral sclerosis (Pandarinath et al., 2015), and so these results demonstrate the ability to isolate units beyond the motor cortices also.

It is worth noting that Barrese et al. (2013) described a period of quiescence for a period of weeks immediately post-implantation, prior to the appearance of units that were persistent for many months, and it was beyond the scope of this study to assess whether the isolatable units found during these initial days and weeks persist into the longer term. As such, it is plausible that the stable units found over a matter of days here would actually be lost in the transition from subacute to chronic recordings, from a protracted response to the initial implantation of the device over the first weeks. On the other hand, in the 30 minute, in depth analyses, the highest SNR, the highest mean detection multiple, and the lowest probability of false negatives were all found in patient 2, and it was in this patient that the

array remained implanted for the longest. The epoch analysed for this patient was from over 3 weeks post implantation rather than a matter of days in the other 3 patients, and so this simultaneously provides anecdotal evidence in support of the slow appearance of the highest quality single unit recordings.

It is important to note that the MEAs were necessarily implanted into putative ictogenic regions, and so the ability to sort spikes into single units may be impeded by pathological alterations to the cortical tissue. None of the patients studied here had cortical dysplasia or apparent lesions (Table 2.1), however, microscopic pathological alterations may still occur in tissue that is frequently recruited to the seizure. For example, post-resection analysis of neocortical tissue from patients with temporal lobe epilepsy has found an increase in mean volume of the somata, sparser dendritic trees, and reduction of dendritic spine density (Bothwell et al., 2001).

While it is plausible that such alterations may hinder spike sorting, it should be considered that, first, the larger volume cells in cortical dysplasia are also accompanied by larger voltage-gated currents (Farrell et al., 2008), and thus, if anything, should show extracellular spikes of increased amplitude, with the caveat that the quiescence of "balloon cells" may reduce the probabilities of the electrode tips being near an active neuron. Second, the most single units successfully isolated per patient in this study were in patients 1 and 2 (Table 3.2), in both of whom the MEA was implanted within the territory fully recruited to the seizure (see Chapter 4). While control studies incorporating equivalent MEA recordings in patients without epilepsy (which are ethically unfeasible) would be necessary to definitively prove whether spike sorting is affected negatively by alterations to tissue in epilepsy, the successful clustering of waveforms into single units in these patients provides evidence that this is not the case. Further studies with reference to cases of cortical dysplasia, controlled to cases without evident lesions are necessary to confirm this observation, however.

### 3.3.1. *Stability of electrode features and single unit characteristics*

The methods of comparing the probability densities for the electrode features as shown in Figures 3.1 – 3.3 necessitates using only the first two principal component scores for each detected waveform, thus removing some of the features upon which the waveforms may have been separated, since it has been shown that 4 or more principal component scores are generally optimal (Adamos et al., 2008). Even so, the difference between the intra- and

inter-electrode correlation coefficients is evidence that, even using only the first 2 principal components scores, there was enough dissimilarity between single unit waveforms such that electrodes had measurably different characteristics.

These differences between the intra- and inter-electrode correlations are important, since no definitive knowledge is available for whether a waveform from one epoch arises from the same cell as a waveform from the next epoch. The majority of extracellular action potentials will be unipolar, negative deflections resulting from $Na^+$ ions flowing away from the electrode tip into the cell (note that the after-hyperpolarization is only evident in the extracellular record if the recording electrode is very close to the cell body (Gold et al., 2006)). As a result, waveform similarity alone is not a good indicator of two spikes having arisen from the same neuron. To counter this, in experimental set ups in which it is plausible to progress the electrode array slowly through the cortex, it is possible to derive a null distribution of waveforms that unequivocally arise from different units (Tolias et al., 2007). This distribution provides a measurement for the expected level of matching units arising purely by chance.

In these human recordings, however, slowly moving the electrode tips to provide the null distribution is unfeasible. Instead, we can use recordings from different electrodes, since the smallest inter-electrode distance (400 μm) is greater than the listening sphere of the electrodes for the voltages required to record single unit spikes. As such, the continued significant difference between the intra- and inter-electrode correlation coefficients shown in Figure 3.3 A, up to nearly 48 hours after the comparison epoch, provides strong evidence that even while anticipating drift in the underlying wave shapes, the cleanly defined single unit clusters from one day are likely to correspond to those in the following day, and likely beyond. On average 36%, and up to more than 50%, of this reduction in correlation coefficient over the 48 hour period is explained by allowing for a drift in the total features in the principal component space through the use of cross correlations. This provides evidence that much confidence can be gained in future long-term recordings through the storage of brief epochs during non-experimental periods, in situations where continuous data storage is not an option.

The discrepancy between the correlation coefficient and the maximal cross correlation coefficient implies strongly that major alteration in the underlying wave shapes over time is often a drift of the shape of the extracellular action potential recorded at that electrode.

Thus, storing even brief epochs every hour allows for an assessment of the likelihood that a discrepant waveform corresponds to the same unit, without major increase in data storage, or the need to separately spike sort all epochs between experimental sessions. This comes with the caveat that the principal components must be calculated on all epochs simultaneously. This was important to be able to directly compare the scores for the waveforms from one epoch to another, though in doing so certain features which may have resulted in distinct clusters in any single epoch may have been lost, because these are non-stationary data resulting in clusters becoming smeared over time. In other words, calculating the principal component features on separate epochs independently may well provide the ability to isolate more units, but in doing so, the ability to interpret the resultant activity across epochs is lost. In the event that continuous data may be analysed over such a long period, given the necessary computational power, utilizing a non-parametric Bayesian approach to spike sorting would likely provide optimal results, as these would allow tracing of units through any drift in principal component scores over time (Bar-Hillel et al., 2006; Wood & Black, 2008).

The use of the aforementioned electrophysiological signatures may aid in mapping slow alterations in waveform; a situation that is likely from micro-motion of the MEA, or of microscopic movements of cells relative to the electrodes. However, I also show evidence for sudden alterations in the extracellular action potential recorded from a single cell in Figure 3.9. In these instances, I have shown through a continuous analysis of multiple 30 minute epochs, that the SNR, amplitude and firing patterns of single units are stable. Therefore, in the case of any sudden alteration of waveform, a detailed look at the full feature set of putative single units through time will allow for easy assessment of which alterations correspond to movements of a single cell, as opposed to the loss of one and subsequent gain of another cell at that electrode. Further to the above metrics, these results also confirm that even in the event of significant drift from one day to another in feature space, the intrinsic firing properties of single cells are a strong indicator of their identities, from one day to another.

### 3.3.2. *Putative sub-classification by cell type*

Finally, I attempted a sub-classification of these units through a subset of the above metrics, along with features from the average wideband signal extracted from when the spikes from

a single cell are deemed to have occurred. In order to garner the most information from these single unit recordings in any situation, the putative single neurons must be separated by cell type, so as to be able to assess network properties with knowledge of the output of that cell: whether it inhibits or excites downstream neurons. Indeed, a key direction for these studies would be to elucidate the spatiotemporal dynamics of these subpopulations of neurons at various points of epileptiform activity: inter-ictal; pre-ictal; ictal; and post-ictal – and learn how these time points differ in their inhibitory network dynamics in order to give rise to these epileptic events. An accurate sub-classification of these cell types would be able to give significant insights into at what point the feed-forward inhibitory surround gives way to allow propagation of the hyper-synchronous activity during seizures.

My initial attempts at sub-classifying in this manner indicated that no individual measurement alone provided enough information on which to separate the two populations, throughout all patients. Indeed, in two of the four patients, no clear ability to sub-classify on a mixture of these measurements was evident, and relying on a *k*-means algorithm in this situation is unwise, as its nature is to separate the dataset into the specified number of clusters, regardless of how statistically likely those clusters are. That is to say, the *k*-means algorithm will provide the resultant *k* clusters that are most likely such that membership to a cluster will be the one with the closest centroid, though it makes no assessment as to the reasonability of applying a separation between data points at the boundary between clusters; a perfectly homogeneous dataset will still be clustered, but the boundaries between the resultant clusters will be no less dense than that of the centre of each cluster.

In the two patients where a sub-classification was attempted, and in the equivalent results from a global sub-classification, the firing patterns of the two putative populations were not conclusively indicative of a successful separation. However, the trend was indeed towards the units with briefer action potentials (the fast spiking units) having autocorrelations indicative of the activity found in confirmed fast spiking interneurons (Csicsvari et al., 1998). As such, while the ability to conclusively separate the two populations is unlikely in this dataset, the correlation between the briefer action potential in fast spiking interneurons due to the fast and intense activation of $K^+$ channels shortly after the $Na^+$ channels, and the firing activity of these cells is evident in these data also.

Why then, if this relationship appears to be true also in these recordings, is it not feasible to sub-classify the units as it has been shown in animal studies (Barthó et al., 2004; Csicsvari et

al., 1998; 1999; Mitchell et al., 2007)? The majority of spike sorting studies has been done on recordings from the pyramidal cell layer of the hippocampus. This layer is unique in its densely packed, and well ordered neurons, of only two significantly differing cell types, in stark contrast to the neocortex (Shepherd, 2004). It is likely that the ordered organization of the hippocampus results in cleanly decipherable wave shapes from the two populations due to the relative position of the soma to the electrode tip being a primary source of waveform variation (Gold et al., 2006). Furthermore, antidromically identified pyramidal cells in the motor cortex have been shown to produce waveforms with kinetics previously thought to characterize fast spiking interneurons (Vigneswaran et al., 2011), a situation which may be unique to neocortex. However, this might only be the case in motor cortex due to the large Betz cells found exclusively in M1, and thus not an issue in other regions of neocortex, and in fact in the first of the aforementioned animal studies, this separation was performed in neocortical recordings from rats.

A major contributing factor to these successes may be in sheer volume of single units, and in the accuracy in spike sorting afforded by high enough density recordings to pick up activity from a single neuron on multiple electrodes, thereby providing further isolation of units based on source separation (Takahashi et al., 2003). The higher density of such electrode arrays also increases the likelihood of recording from single units that truly have monosynaptic connections to one another – over a distance of at minimum 400 μm between electrodes, true monosynaptic connections are unlikely. This was confirmed, in part, through the finding that any putative monosynaptic connections in this dataset were found to be between units from the same electrode as one another. Notably, though, the number of spikes found in these putative connections is low, at only 7, 5 and 14 each (Figure 3.15). As a result, these connections could likely be deemed to arise purely due to low sampling issues, causing limited coincidences able to surpass the confidence intervals arising from the 1,000 trials. That said, in that instance, more connections would likely be found amongst separate electrodes by chance alone, and none were found here. Furthermore, despite these low numbers, the percentage of spikes from the presynaptic cell that had corresponding spikes in the postsynaptic cell (5.20%, 7.69% and 2.40%), are not insubstantial. Though these initially appear low, the transmission rate from one excitatory cell to another might be expected to be lower, given the usual need for many coincident excitatory postsynaptic potentials in order to elicit an action potential. A single excitatory postsynaptic potential

corresponding to a resultant action potential at the probabilities found here would allow for the complexities required for cortical processing; if neurons were very likely to consistently cause action potentials in the neurons that they innervate, the potential for information latent amongst the network would be considerably decreased. That said, this is not a necessity in all cases; in the instance of axo-axonic cells, the post-synaptic effects of single neurons are capable of eliciting action potentials, monosynaptically (Szabadics et al., 2006).

Although the putative monosynaptic connections found here were limited to units from the same electrode as one another, a recent study from Peyrache et al. (2012), has shown sub-classification of cell types, and confirmed them with unit interactions of a supposed monosynaptic nature across electrodes, in human neocortex from equivalent recordings to those performed here. How might these cross correlations come about over such a large distance? I would posit that given the concurrent input across the large network, coincident spiking, or lack of activity, within a brief window of ± 4 ms as described, is actually more likely than would be expected based on the 99% confidence from the number of spikes alone, as the entire network of neurons is relatively entrained to the underlying oscillations. These coincidences as such would not be reliably indicative of monosynaptic connections, but rather of a large number of cells firing with similar patterns due to the shared inputs to the local network. Regardless of this caveat to the unit cross correlations, the sub-classification performed in their study was convincing based on firing rates, autocorrelations, and waveforms. It is of note that the only obvious discrepancies between the recordings performed by Peyrache et al. (2012) and those described herein, are their use of 1.5 mm shank MEAs as opposed to 1 mm, and the distance from the ictogenic region. It is therefore possible that deeper layers afforded better geometric localization by cell type to allow for sub-classification, or that the proximity to pathological cortex of the recordings analysed here resulted in altered underlying properties of specific cell types.

# Chapter 4. Loss of single unit features in ictal core

## 4.1. Introduction

In Chapter 3 I showed the ability to follow activity from putative single neuronal sources stably, over many hours, during non-seizure epochs, and so I next asked how the occurrence of epileptiform activity might disrupt such ability. Spike sorting analyses of extracellular recordings rely on units showing a stable action potential waveform, and so the occurrence of one or many features associated with ictal activity may confound such studies.

A key consideration is that intense activity at the population level can lead to increases in extracellular concentrations of $K^+$ ions ($[K^+]_o$; Somjen & Giacchino, 1985). Since the duration of an action potential is primarily related to the kinetics of the return current achieved through the opening of $K^+$ channels, increases in $[K^+]_o$ will result in an increased width of the action potential related extracellular spike. Furthermore, with a large increase in $[K^+]_o$ it is possible for neurons to go into a state known as depolarization block. The resultant extreme membrane potential depolarization in this state leaves the cells incapable of firing action potentials due to the deactivation of voltage-gated $Na^+$ channels (Somjen, 2004). Moreover, a computational model has described how the transition from physiological activity, to ictal activity, through to spreading depression can be accounted for through progressive increases in $[K^+]_o$ (Wei et al., 2014). As such, it is likely that cortex that is recruited to the seizure will show increases to $[K^+]_o$, resulting in alterations to spike shape that would impede the ability to spike sort reliably.

Similarly, paroxysmal depolarising shifts (PDSs) are regarded as the intracellular correlate of epileptic discharges, and are characterised by reduced amplitude action potentials riding on the crest of the large depolarisations (Traub & Wong, 1982). However, PDSs have been described almost solely from animal studies, where intracellular recordings are feasible. Their occurrence in human epilepsy remains controversial, with several studies of human extracellular recordings failing to find any evidence of PDSs (Babb et al., 1987; Stead et al.,

2010; Truccolo et al., 2011; Wyler et al., 1982). One explanation for these null findings relates to the location of the electrode placements.

The terms “ictal core territory” and “ictal penumbra” have been proposed to distinguish between two qualitatively different activity patterns during seizures (Schevon et al., 2012; Trevelyan & Schevon, 2013; Weiss et al., 2015). These concepts have arisen from a large body of animal studies describing epileptiform activity in terms of localised, intense, hypersynchronous discharges (ictal core), with many neurons displaying PDSs (Kandel & Spencer, 1961a; Matsumoto & Marsan, 1964; Traub & Wong, 1982), associated with a marked inhibitory response in surrounding territories (penumbra) (Cammarota et al., 2013; Dichter & Spencer, 1969a; 1969b; Prince & Wilder, 1967; Schwartz & Bonhoeffer, 2001; Timofeev et al., 2002; Timofeev & Steriade, 2004; Trevelyan et al., 2006; Trevelyan et al., 2007; Wong & Prince, 1990), which appears to restrain the propagation of the focal pathophysiology. Note however, that even in the penumbra, there is typically increased firing over baseline activity levels (Schevon et al., 2012). Recent studies of microelectrode array recordings in human subjects have now extended this description to include spontaneous (“habitual”) human seizures. In both core and penumbral territories, one can find very large, low frequency local field potentials recorded on subdural recordings from macroelectrodes on the surface of the brain. Of critical importance, however, is that these recordings also allow the two territories to be distinguished by means of a pronounced phase-locked high gamma band signal visible only in the ictal core (Weiss et al., 2013).

It is important to be clear about the distinction between the widely used term, the “seizure onset zone” and the ictal core. The former is the best estimate of where a seizure begins, and is a fixed location reflecting what is happening at a single time point, namely the earliest time when seizure activity could be detected. In contrast, the ictal core is a dynamic concept, describing which cortical territories are fully recruited at a given point in time, and so changes over the time course of the seizure.

In this chapter, I present analyses of human single neuronal unit activity from both penumbral and core recruited territories (Merricks et al., 2015). I asked if the ability to follow single units, as described recently by Truccolo et al. (2011), might depend on whether the MEA was located within the core or the penumbral territories at a given time point during a seizure. I therefore investigated the stability of these units during a continuous epoch, from half an hour immediately prior to seizure onset, through the ictal activity, and

immediately post seizure termination. I assessed how unit activity was altered during the seizure, and documented the recovery of unit activity after seizure termination.

## 4.2. Results

### 4.2.1. *Intense firing is associated with changes in extracellular spike shape*

The PDS is regarded as the intracellular hallmark of an epileptic discharge. There is a strong presumption that these changes should also be apparent in extracellular recordings, but to check this explicitly, I analysed previously performed cell attached recordings, provided courtesy of Dr Andrew Trevelyan. These recordings were made from visually identified, layer 5, pyramidal neurons in brain slices prepared from young mice, and were analysed during epileptiform events that had been induced through the removal of $Mg^{2+}$ ions from the bathing medium (0 $Mg^{2+}$ model; Figure 4.1). Cell attached recordings isolate the membrane currents of the patched cells without interference from activity in neighbouring neurons. These recordings showed that during the seizure-like event, there occurs a large decrease in the amplitude of the action potential currents, with an associated broadening of the spike (Figure 4.1 C). The scatter plots of the amplitude of the action potential current versus the inter spike intervals (ISIs) showed several notable features (Figure 4.1 D).

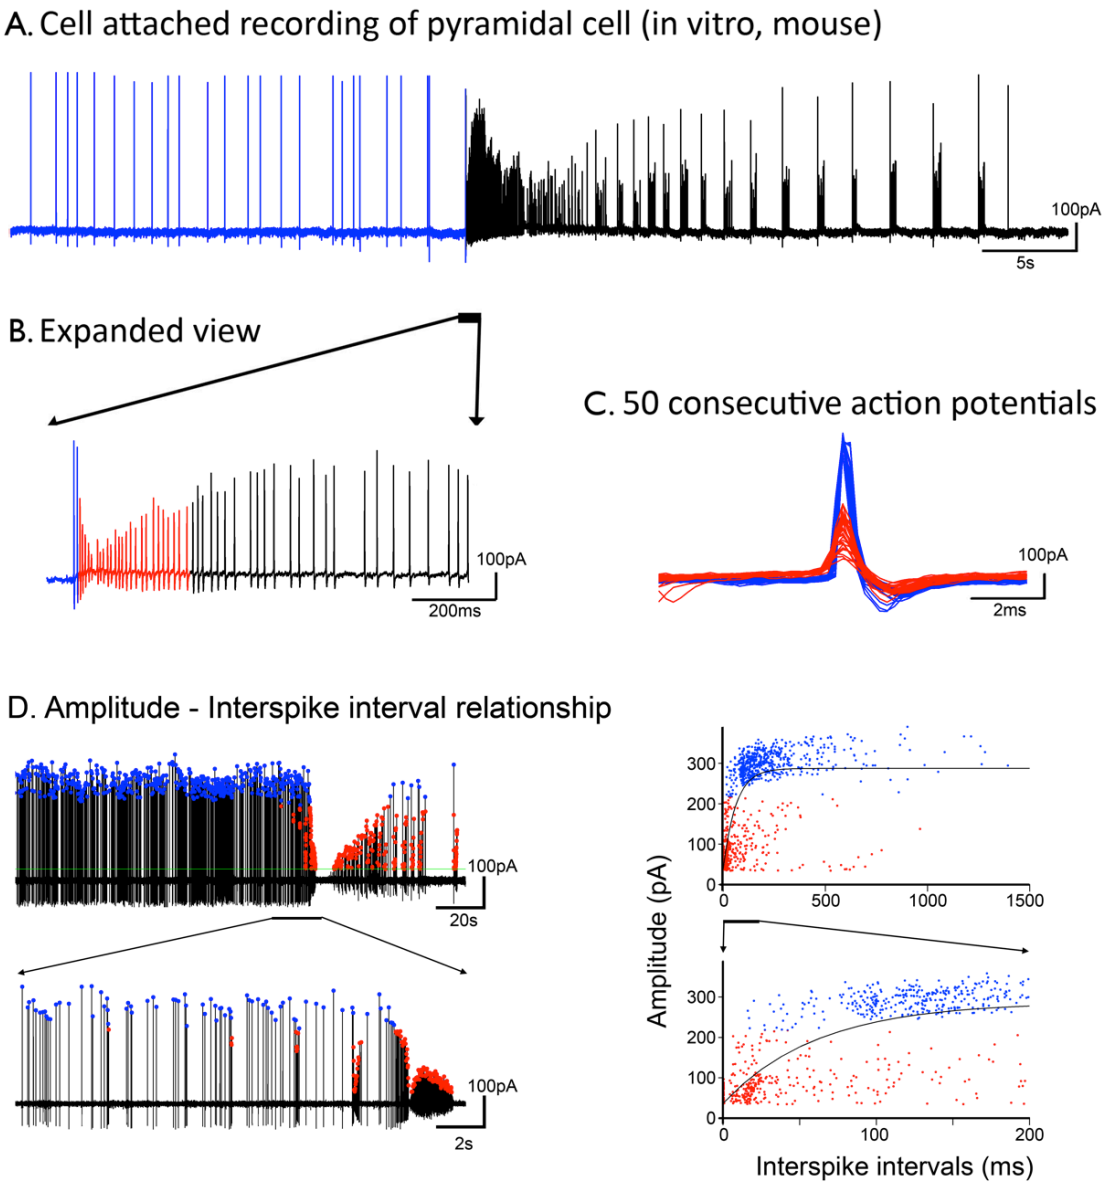


**Figure 4.1. Loss of confirmed single cell extracellular waveform during epileptiform activity in an animal model.**

**A.** Cell attached recording of a pyramidal cell in an *in vitro* model of epilepsy (0 $Mg^{2+}$) showing the shortening and broadening of waveform characteristic of the paroxysmal depolarizing shift due to the failure to deinactivate the $Na^{+}$ channels. **B.** Expanded view of the first second of epileptiform activity showing severe alteration in spike shape **C.** 50 consecutive action potentials (25 pre-recruitment (red), 25 post onset (blue)) overlaid, showing the large change in shape. **D.** Cell attached recording with the times and amplitudes of peak action potential currents marked (automated detection above a threshold (green line) of 10 times the median absolute deviation divided by 0.6745; see Quiroga *et al.* (2004)). An arbitrary amplitude threshold was used to colour code the smaller events red, and the larger ones blue; this simple distinction separates out the spikes during the epileptiform event from those in the baseline period. Note particularly the large and overlapping range of inter-spike intervals for both subsets of action potentials, in the right panels. A single exponential fit is shown for the entire data set ($\tau$ = 75.9 ms).

First, there was a clear evidence for a refractory period in all cells, with the shortest ISIs being followed by the smallest action potentials. Exponential fits to the amplitude versus ISI plots gave a time constant, $\tau$ = 108.9 ± 59.6 ms (n = 5 cells from 5 animals). Second, by identifying the smallest amplitude action potential current during the baseline period, and setting a threshold just below that, one could neatly separate two groups of action potentials: those occurring during baseline (large) and those occurring during the seizure (small). The most intense firing rates (smallest ISIs) occurred during ictal events, but the wide and overlapping range of ISIs for both the small and large action potentials showed the seizure activity was associated with a greatly extended refractory period relative to the baseline activity. The likely cause of these alterations to firing rate and action potential amplitude is the huge level of synaptic drive during these *in vitro* epileptiform events. Large synaptic currents are also apparent during human seizures in vivo, and so warrant a re-examination of the human spike sorting data during epileptiform activity.

### 4.2.2. *Single unit identification during pre-ictal period*

I therefore performed spike-sorting analyses, as described in section 2.2 and previously performed in chapter 3, in order to identify putative single neuronal units during an epoch from 30 minutes prior to seizure onset, until 15 minutes post seizure termination, where data were available, on microelectrode array (MEA) recordings in four patients with intractable focal neocortical epilepsy. I was able to isolate 305 putative single units successfully from a subset of channels in all patients during these epochs (Table 4.1), using a standard negative voltage thresholding technique and clustering of principal component features. Post hoc accuracy metrics as described in section 2.2.5 for the single units during these epochs provided a false positive estimate of 0.204 ± 0.769% (mean ± 2 S.D.), calculated from the proportion of spikes within unit that violated the 2 ms refractory period. The false negative estimate, made from the proportion of a Gaussian fit to the distribution of detected thresholds which lay sub threshold, was 0.003 ± 0.059% (mean ± 2 S.D.); both metrics as described in Hill et al. (2011).

| | | PRE-ICTAL | ICTAL | | POST-ICTAL | |
|---|---|---|---|---|---|---|
| | | Units (chans) | Units (chans) | Maintained | Units (chans) | Maintained |
| CORE | Pt 1 | 59 (35) | 0 (0) | 0% | 57 (35) | 97% |
| | Pt 2 | 110 (49) | 0 (0) | 0% | - | - |
| PENUMBRA | Pt 3 | 61 (30) | 58 (30) | 95% | 61 (30) | 100% |
| | Pt 4 | 75 (32) | 57 (27) | 76% | 68 (32) | 91% |

**Table 4.1. Single units isolated during different pathological activity with respect to location**

The plot of the first two principal component scores for the collected units from a given electrode has a distinctive distribution, a kind of electrophysiological signature, and so was used to assess stability of the recordings during the peri-seizure time points, as per section 3.2.1. Because many such signatures were being collected simultaneously from multiple electrodes, we can use this between-electrode variance to assess the temporal stability of the recordings, as was applied in Chapter 3 during the non-contiguous, long-term analyses. The between-electrode comparisons are by definition comparisons of different units, and so provide us with an estimate of the variance of the electrophysiological signatures.

To do this, I smoothed the probability density histograms of the principal component scores (normalized in feature space, Figures 4.2 C & 4.5 C) of contiguous 60 second epochs, from 10 minutes prior to seizure onset, to 5 minutes post. Cross correlations were first performed on these smoothed plots relative to a baseline-equivalent, smoothed, probability density histogram from a 60 second epoch from 30 minutes prior to seizure onset (Figure 4.6 A & B, red). These cross correlations were performed on all channels that showed evidence of activity distinct from the homogeneous background multi-unit activity during this epoch (patients 1-4: n = 35; 49; 30; 32 channels respectively).

Cross correlations were then also performed between all electrodes within each epoch, for each patient separately (Figure 4.6 A & B, blue). This demonstrated a low similarity of different electrode recordings seen during baseline periods, a result that was highly reproducible between patients (global mean ± 2 S.D. correlation coefficient from non-seizure epochs for all patients: 0.14 ± 0.01), and which represents a large spread of possible electrode signatures. In contrast, the within electrode, temporal analyses showed extremely high correlations for all non-seizure time points, relative to the baseline epoch taken 30 minutes prior to seizure onset, and again was highly reproducible between patients (global

mean ± 2 S.D. correlation coefficient from non-seizure epochs for all patients: 0.97 ± 0.004). The activity from single units during the pre-ictal period was therefore continuously stable, beyond the level that would be expected by chance due to the similarity between extracellular action potentials from different neuronal sources.

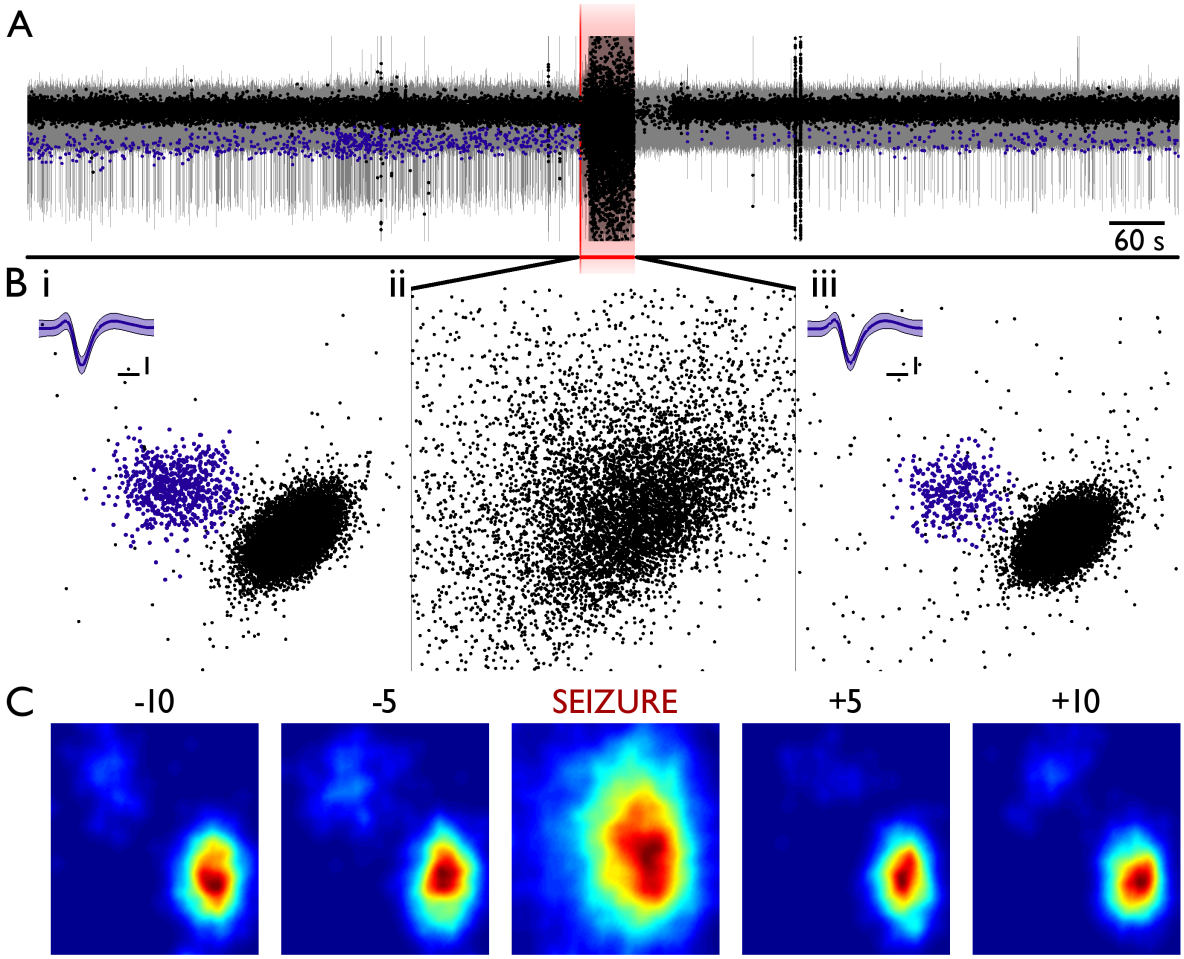


**Figure 4.2. Loss of unit-specific features during seizure in core, recruited territories.**
**A.** Example trace from one channel in Patient 1 (recruited territory) for 10 min prior to seizure onset, to 10 min post-seizure termination. The first principal component values from detected spikes have been overlaid; the blue data points all represent spikes from a single, cleanly separated unit. Note the stability of the principal component scores up until seizure onset whereupon all discernable features are lost. The same unit as in the preceding 10 min can also be seen to make a recovery in the following 10 min, albeit with a much diminished firing rate. **B.** First two principal component scores are plotted against each other. They show a clearly separable unit in blue and, inset, mean waveform (± 2 SD) from: (**i**) 10 min prior to seizure; (**ii**) during seizure (60 s total epoch); and (**iii**) 10 min following seizure termination. Note the maintained waveform and corresponding principal component scores in (iii) relative to (i), despite no evidence for the same unit during the seizure in (ii). Both axes are maintained throughout, and principal components were calculated simultaneously. Waveform inset scale = 0.2 ms, 20 µV. **C.** Kernel density estimate histograms (sigma = 5, 100 bins) of principal components 1 (abscissa) and 2 (ordinate), from 60 s epochs from 10 and 5 min prior to seizure onset, during the seizure, and 5 and 10 min post-seizure termination, normalized to the dimensions found in the 10 min prior epoch. Red denotes the highest probability of finding a spike while blue denotes the lowest.

### 4.2.3. *Spike sorting differences between the penumbral and ictal core territories*

I next analysed the stability of unit spike shape around the times of seizures relative to the location of the MEA. The positioning of the MEAs was made on the available clinical evidence about the best estimate of the location of the seizure onset zone. Subsequent analyses showed that in two patients, this territory was fully incorporated into the core recruited territory during seizures (Patients 1 and 2), while in the other two patients the MEA was judged to have been placed in the penumbral territory (Patients 3 and 4) (Schevon et al., 2012). This important point is worth restating: that even within the clinically defined seizure onset zone, some territories were not fully recruited to the ictal core territory, illustrating the difficulties of defining the seizure onset zone precisely. I made use of this chance distinction between the recording locations, to determine the stability of the unit recordings over time, with particular interest in whether single units could be followed during seizures in both the penumbral and the ictal core territories.

In recordings from MEAs that eventually were fully incorporated into the ictal core territories, there were also preceding periods of penumbral activity, as defined by the large amplitude low frequency signal with low levels of spiking. Notably, for these "penumbral periods", spike shapes were preserved in all electrodes (Figure 4.3). At the arrival of the ictal wavefront, however, and throughout the seizure subsequently, many action potential waveform characteristics change, including the principal components, wavelet features, waveform maxima and minima, and the non-linear energy (Figures 4.2 - 4.4). In short, all features on which spike sorting might be performed are altered, preventing any useful strategy for following a single neuron's activity patterns through the seizure.

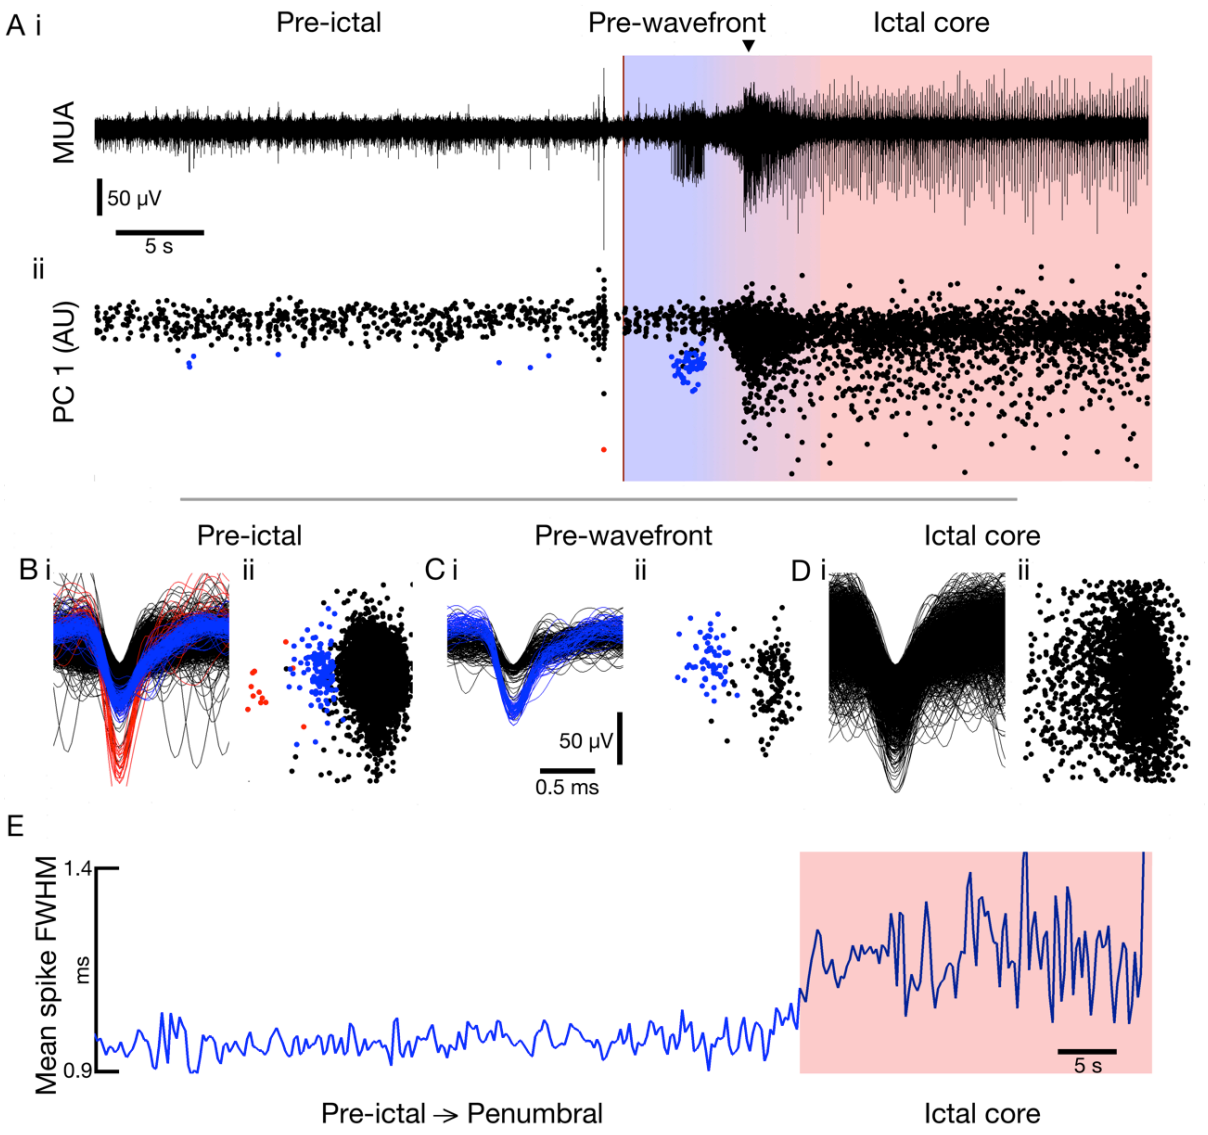


**Figure 4.3. Maintained waveform during propagation prior to local ictal onset.**
Cortical regions that are later incorporated into the ictal core show maintained waveforms after seizure onset while they are still in the penumbra (prior to the wave front reaching and incorporating that territory). **A.** (**i**) Example MUA (300 to 3000 Hz) trace from Patient 2 and (**ii**) accompanying first principal component through time of detected spikes, from 30 s prior to global seizure onset (dark red line) until 30 s after. For the first 4.8 s after seizure onset the territory local to the electrode remains penumbral, prior to the wave front reaching the region (blue shaded area), after which the territory becomes incorporated into the ictal core (red shaded area). Note the maintained spike height and first principal component scores while in the penumbra, and loss of specificity once incorporated into the ictal core. **B – D.** (**i**) Waveforms and (**ii**) their first two principal components from the channel shown in A, from 5 min pre-ictal, the penumbral 4.8 s, and the ictal core, respectively. Note the maintained waveform and associated features of the blue unit during the penumbral period (C) followed by obscurement when incorporated into the ictal core (D). **E.** Mean spike full-width at half-maximum (FWHM) of all single units from Patient 2, aligned in time to local incorporation into seizure (ictal core shown in red shaded area), showing that even when including spikes from the same principal component region despite a lack of defined clusters, the spike width within clusters increases at local onset, and not while in the penumbral region.

After recruitment there was a notable increase in the number of large amplitude 1st principal components (Figure 4.2 B), suggestive of superimposed spikes caused by hypersynchrony, and reflecting previously observed synchronized bursting across all electrodes in the 4 x 4 mm MEA at this time (Weiss et al., 2013). When these features are plotted out in their respective feature spaces, the clusters that define a single unit get lost inside a great cloud of data points for all feature types. I therefore considered two alternative hypotheses to explain the inability to cluster. First is the occlusion hypothesis, which is that units preserved their shape, but clustering was prevented by the appearance of other units. The second is the PDS hypothesis, which is that spikes change their shape. To distinguish between these two hypotheses, I plotted the change in spike half width for units throughout the seizure, from within the same small area of the principal component space as the original cluster (Figure 4.3 C & D). This showed an increase in spike half width at the time of recruitment for data points within that area of the feature space, thus indicating that the occlusion hypothesis alone cannot account for all observations, but remaining consistent with the PDS hypothesis.

The within electrode temporal correlations were high for all inter-ictal epoch comparisons (referenced to a time point 30mins before the seizure; Figure 4.6 A & B, red bars). Within the core territory (Figure 4.6 A), this correlation dropped significantly during the seizure compared with the baseline reference epoch (patient 1, seizure 1: 47% drop from baseline, $p << 0.001$, n = 35; patient 2, seizure 1: 10.4% drop, $p << 0.001$, n = 49; Mann-Whitney U test). There were instances during inter-ictal periods when the temporal correlations dropped significantly, but this was always a far smaller change than shown during the seizure event (patient 1, maximal inter-ictal shift = 1.3%, p = 0.02; patient 2, maximal inter-ictal shift = 0.8%, n.s). The ictal core activity was also characterised by a significant rise in the between-electrode cross correlations (Figure 4.6 A, blue bars; comparison of seizure between-electrode distribution with the equivalent between-electrode distribution from the epoch 10 minutes prior to seizure onset; patient 1: $p << 0.001$, patient 2: $p = 0.003$, Mann-Whitney U test), which occurred because the electrode signatures lost their most distinguishing features (i.e. the discrete clusters which represent single units). Note also that such alterations were not found in either the intra- or inter-electrode measurements in the long-term equivalent analyses (c.f. Figure 3.3 B).

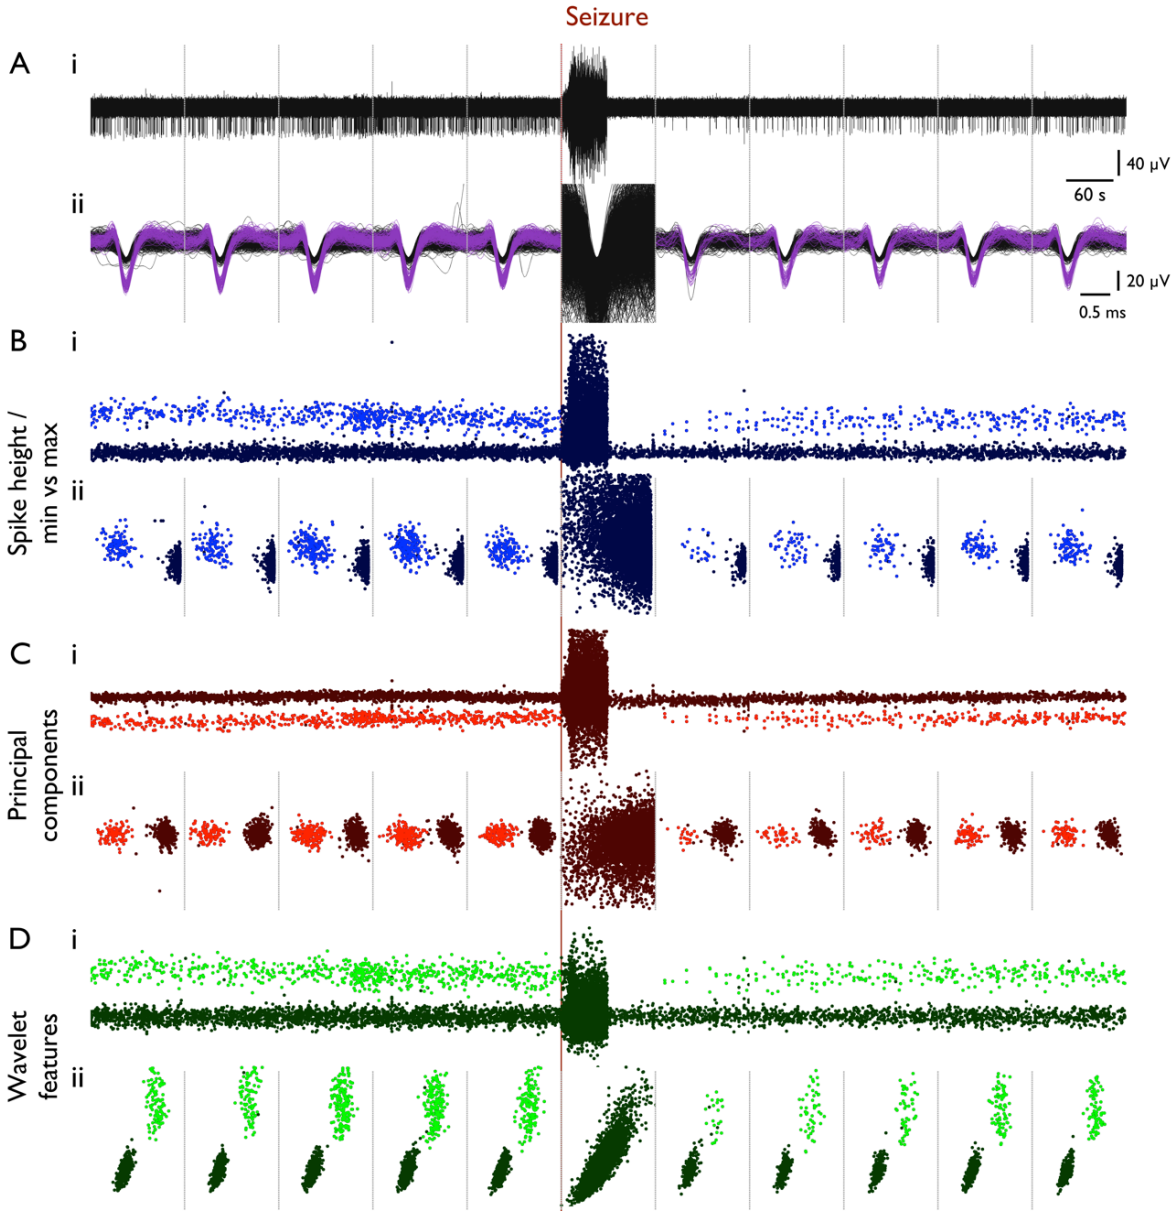


**Figure 4.4. Loss of specificity in core territory in all cluster-able features.**
**A.** (**i**) Example trace from one channel in Patient 1 (recruited territory) from 10 min prior to seizure onset until 12 min post-seizure onset, segmented into 2-min epochs (grey dotted lines, used throughout). (**ii**) All waveforms from each 2-min epoch shown in **i**, with background noise in black and single unit in purple. **B – D.** Multiple cluster-able features, with single unit as lighter colour throughout. **B.** (**i**) Spike height plotted versus spike time and (**ii**) spike minimum versus maximum for each epoch. **C.** (**i**) First principal component versus time and (**ii**) first versus second principal component for each epoch. **D.** (**i**) First wavelet feature versus time; and (**ii**) first versus second wavelet feature for each epoch. Note loss of features throughout during seizure epoch.

In contrast, in penumbral territories, action potential waveform shape and associated features were maintained throughout the seizure for the majority of units, accompanied by preservation of features in principal component space (95% and 76% unit retention during seizures, for patients 3 and 4 respectively, relative to the pre-ictal period, Table 4.1). There was no significant difference in the cross correlation analyses for any time point in the penumbral recordings. This result indicates that in penumbral cortical areas, despite clear evidence for the seizures directly influencing firing rate (Figure 4.5 A), cells did not undergo paroxysmal depolarising shifts, nor was there sufficient superposition of spiking to distort summated waveforms. Because the unit clusters persisted, the between electrode cross correlations also remained very low during the seizures in the penumbra (Figure 4.6 B). This difference between penumbral and ictal core activity was confirmed by a second analytical method based on subtracting the baseline electrode fingerprint from subsequent time epochs, to identify times when there occurred large deviations in the signature (Figure 4.6 C & D).

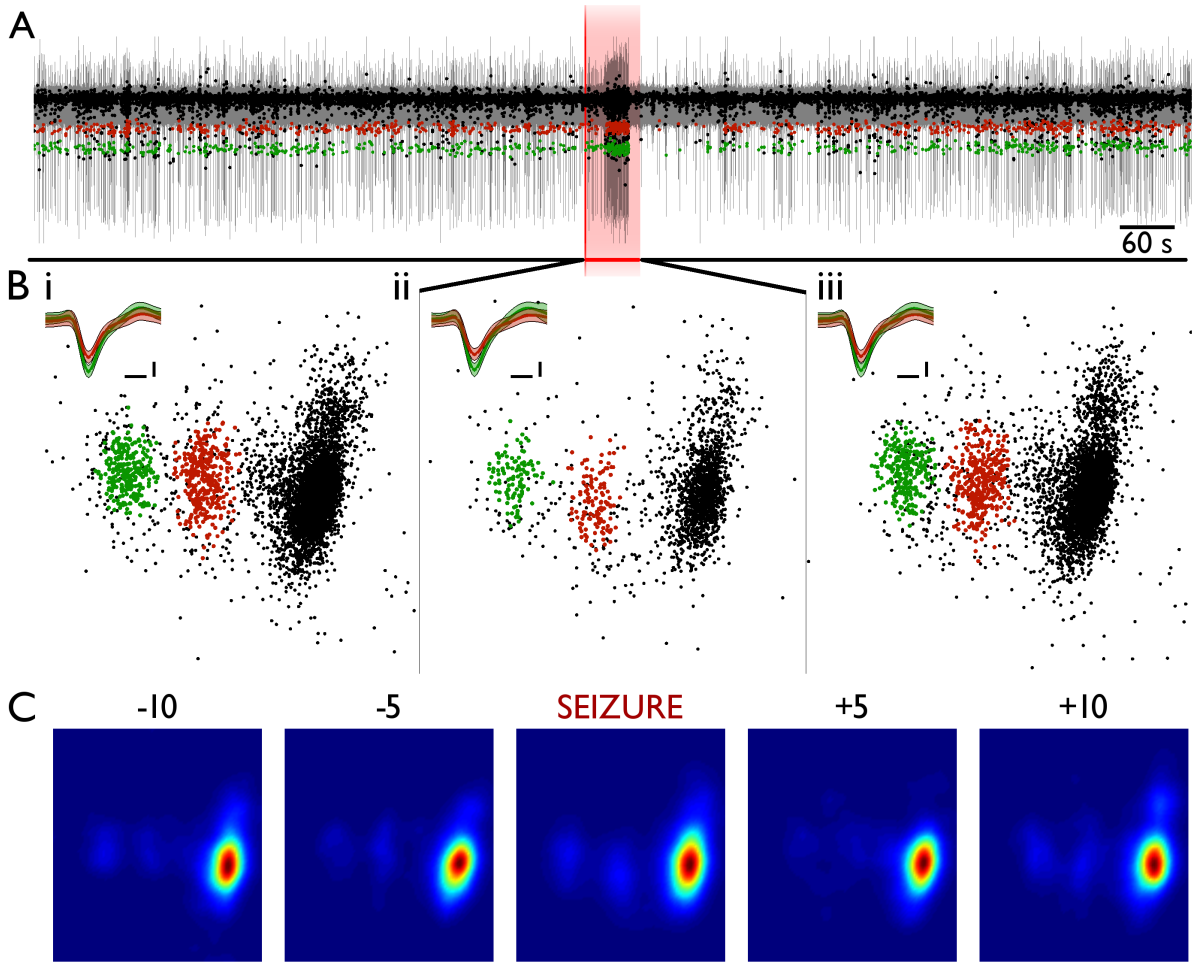


**Figure 4.5. Maintained waveform and features during seizure in penumbra.**
Same format as for Figure 4.2, showing the preserved electrophysiological signatures recorded within the penumbral, non-recruited territory (Patient 3; shown in Figure 7g as "Seizure 1" in Schevon *et al.* (2012)). The large amplitude, low frequency signals indicative of seizure activity, corresponded approximately to the period of increased unit activity. Note that the pink shaded territory indicates the sampling period for B (ii) [60 s; the same duration as in Figure 4.2 B (ii)], but that in this figure it also includes short pre- and post-ictal periods. **A.** The 300 Hz to 3 kHz bandpass filtered trace from a single electrode, with two well separated units indicated in green and red. **B.** The distributions of the first two principal component scores for these units, together with many other units, which are shown in black. All scales are maintained through, inset scale = 0.2 ms, 20 µV. **C.** The Kernel density estimate histograms for the distributions shown in B. Note how the two units are clearly separable during each time point, and that the principal component scores (and Kernel density estimates) are stable, even though there is a marked rise in firing frequency during the seizure relative to periods both preceding and following the seizure. The subdural EEG from this seizure is shown in Schevon *et al.* (2012; Figure 7g, "Seizure 1").

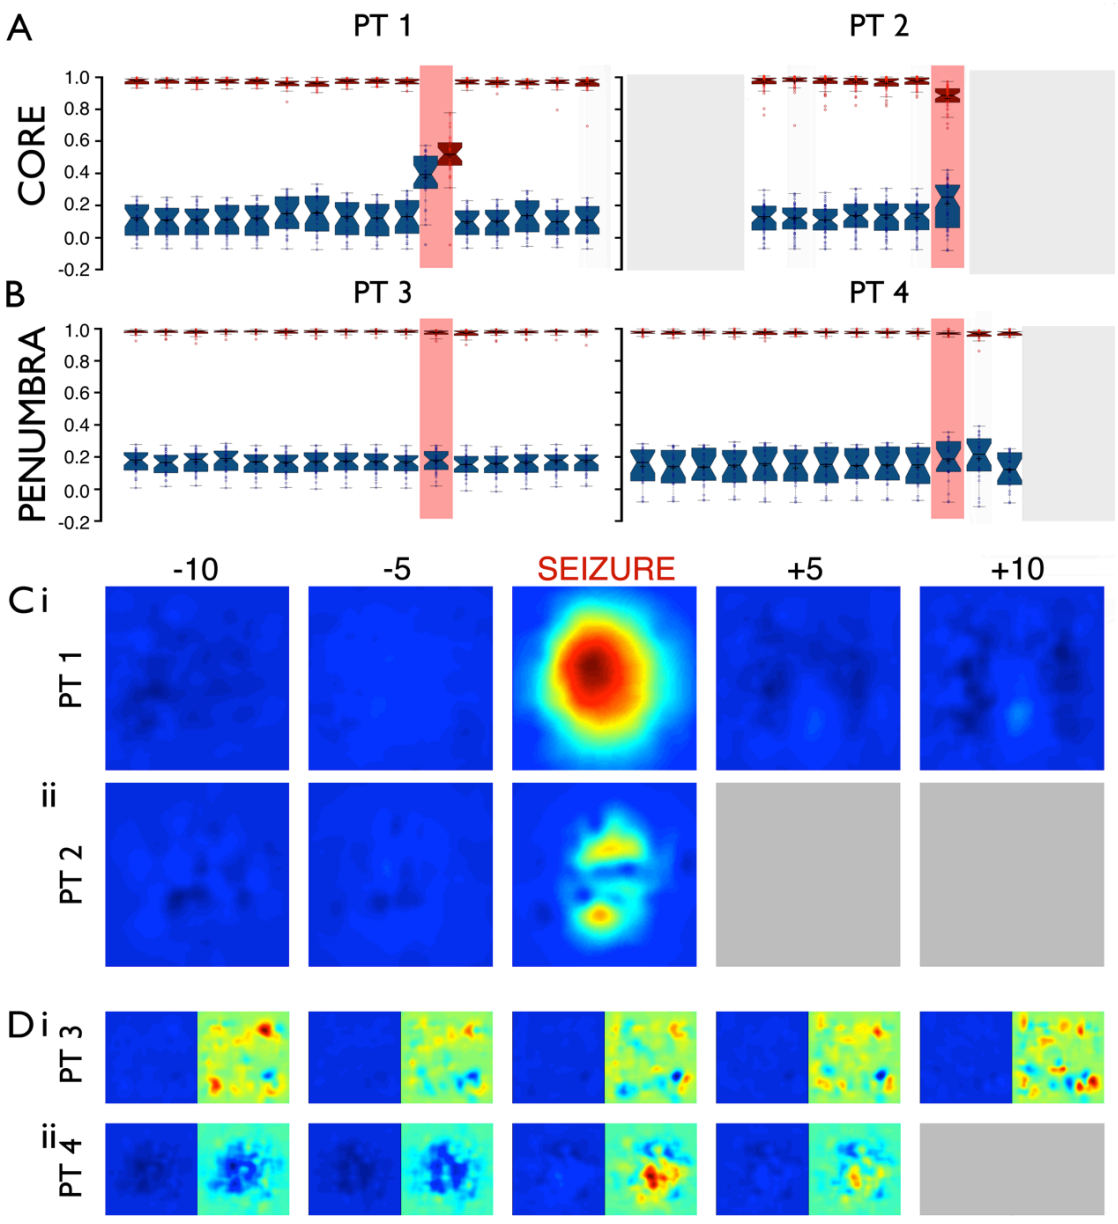


**Figure 4.6. Loss of separable features from single units is characteristic of core recordings, not present in penumbra.**

**A** & **B.** Cross correlations of Kernel density estimate histograms as shown in Figures 4.3 C and 4.5 C, for 60 s epochs from 10 min prior to seizure onset to 5 min post termination in all four patients [core in A: Patient 1 (30 electrodes), Patient 2 (45); penumbra in B: Patient 3 (30), Patient 4 (22)]. Red boxplots show intra-electrode cross correlation coefficients relative to the Kernel density estimate histogram 30 min prior to onset. Blue boxplots show interelectrode average cross correlation coefficients relative to each other electrode during that epoch. The strong correlations over time within each electrode during baseline periods compared to the weak correlations between electrodes show that these features are characteristic of activity specific to each electrode. During the seizure, in core recordings (A), the intra-electrode correlation is lost due to breakdown of features and the similarity between electrodes is raised due to similar loss of specificity, as seen in Figure 4.3 B (ii), in all electrodes. Recordings from penumbral territories (B) show no such alteration during seizure with continued specificity within electrodes. Strong correlations within electrode plausibly arise due to well-maintained background noise of more distal cells' spikes. Grey areas denote missing data. **C** & **D.** The mean results over all electrodes of a subtractive method whereby the first Kernel density estimate histogram was subtracted from the others. The background noise of more distal cells' spikes is maintained thereby removing it during subtraction, leaving only alterations in unit specific activity. Resultant mean subtractions from core patients are shown in C, colour normalized to maximal value in Patient 1, showing a larger alteration during seizure relative to other time points. Penumbral recordings show no such alteration during seizure as shown in D, with values normalized to the colour axis in C on the *left*, and normalized within patient on the *right*.

### 4.2.4. *Template matching of waveforms in the ictal core*

It is plausible that spike shapes were in fact maintained whilst within the ictal core, and the inability to spike sort arose merely due to occlusion of the single unit waveforms by a larger variability in detected spikes, rather than the presence of intrinsic waveform alterations. If this were the case, then utilizing a template matching method of spike sorting on the ictal epoch, as described in section 2.3.2, would allow for the detection of the unaltered waveforms corresponding to the pre-ictal units. Therefore, I performed multiple methods of matching waveforms during the seizure, to templates created from the pre-ictal, stable single units. As described in section 2.3.2, statistics were derived from the template for each unit, and every waveform detected during the seizure epoch was compared to each template from that channel. The waveforms were assigned as a match to that unit using different methods for independent subsequent analysis.

Four different template-matching methods were tested for accuracy in the detection of waveforms from single units. This was done through applying each method to the pre-ictal 30 minute epoch, and comparing the resultant number of spikes assigned to the single units to those assigned to the original units from the principal component clustering from which templates were derived.

The first method of assigning matches to a unit was to include all spikes that were within the range of voltages in the waveforms from that unit during the pre-ictal epoch, at every data point. The second method populated units from the spikes in which the correlation coefficient between the waveform and the template unit's mean waveform was greater than the lowest correlation found in the original spikes from that unit. The third group also used the correlation coefficient, but corresponded to the waveforms that were at least the mean minus two standard deviations of the correlations found in the original spikes from the defined pre-ictal unit. The final method of matches was created from all seizure epoch waveforms that had principal component scores that fell within a convex hull created from the cluster of principal component scores from the pre-ictal single unit, thereby selecting all waveforms that were within the same principal component space as the comparison pre-ictal unit.

| Patient | Seizure | Method | | | |
|---|---|---|---|---|---|
| | | Waveform Range | Convex Hull | X-Corr (mean - 2 SD) | X-Corr (minimum) |
| Patient 1 | Seizure 1 | 102.46% | 128.60% | 265.84% | 768.78% |
| | Seizure 2 | 127.57% | 113.06% | 239.56% | 570.47% |
| | Seizure 3 | 115.38% | 104.77% | 210.66% | 435.20% |
| Patient 2 | Seizure 1 | 187.60% | 131.64% | 483.43% | 1386.90% |
| Mean | | 133.25% | 119.52% | 299.87% | 790.34% |

**Table 4.2. Over-estimation of the number of spikes belonging to a cluster using different template matching methods.**

As can be seen in Table 4.2, both cross correlation methods severely overestimated the number of waveforms that were assigned to each unit, throughout all seizures recorded from the territory subsequently recruited to the ictal core. The waveform range and convex hull methods were more accurate, with an average overestimation of 133.3% and 119.5% respectively, where 100% would equate to the same number of spikes being assigned to the unit as were assigned in the more accurate principal component clustering method. Therefore, subsequent analysis of the ictal waveforms was performed using these two methods, since they resulted in the closest number of spikes being assigned to units as were assigned through the well isolated units from principal component clustering. Note that this does not necessarily equate to accurate unit clustering, and these methods would likely include both false positives and false negatives, but they provide the most accuracy for the post hoc assignment of waves to the predefined single units, and so give a reasonable estimate of the presence of similar waveforms during the ictal activity.

| Patient | Seizure | Pre-ictal | Seizure | | | |
|---|---|---|---|---|---|---|
| | | | Waveform Range | | Convex Hull | |
| Patient 1 | Seizure 1 | 15.89% ± 18.01% | 2.70% ± 3.07% | $p << 0.001$ | 7.61% ± 5.81% | $p = 0.002$ |
| | Seizure 2 | 25.31% ± 17.91% | 4.80% ± 4.34% | $p << 0.001$ | 8.82% ± 6.18% | $p << 0.001$ |
| | Seizure 3 | 17.11% ± 14.05% | 4.85% ± 4.39% | $p << 0.001$ | 7.81% ± 6.00% | $p << 0.001$ |
| Patient 2 | Seizure 1 | 13.64% ± 14.90% | 7.35% ± 7.97% | $p << 0.001$ | 7.43% ± 6.20% | $p << 0.001$ |

**Table 4.3. Percentage of total spikes detected classified as an isolatable unit by different methods.**

I therefore calculated the percentage of detected spikes during the seizure that were classified as having arisen from the defined pre-ictal single units, using these methods.

Relative to the number of detected spikes that were assigned to the defined units during the pre-ictal period, both the waveform range and convex hull methods showed a significant reduction in the proportion of spikes being classified during the seizure (Table 4.3).

One possible explanation for such a decrease in the proportion of spikes being classified could be that some units simply did not fire enough to constitute a definable cluster during the pre-ictal period. This hypothesis was supported by the fact that despite the decrease in the proportion of spikes being classified, the firing rate of these template matched units predominantly increased during the seizure. There was a significant increase of firing rate in 93.6%, 87.9%, 95.5%, and 58.2% of single units during the seizure from patient 1, seizures 1 to 3 and patient 2, seizure 1, respectively using the convex hull method, and in 62.8%, 61.5%, 71.9%, and 37.3% of units using the waveform range method. However, both methods were shown to overestimate the likelihood of a waveform truly corresponding to a single unit during the pre-ictal epoch, and the sheer number of spikes during the seizure would necessarily result in a large number matching the shape of the pre-ictal units by chance. The mean number of spikes being assigned to a unit was directly related to the number of detected spikes during the seizure, dropping over time when normalized to this underlying rate of detection (Figure 4.7 A).

The hypothesis that this increase in firing rate was a result of chance matches from the manifold increase in the number of detected spikes during the seizure was in turn further supported by the fact that the post hoc accuracy metric of false positives, by way of the percentage of ISIs within a unit that violated the refractory period (i.e. a unit with two action potentials within a short enough period of one another as to almost certainly arise from separate neurons due to the absolute refractory period), increased significantly in all seizures using the convex hull method, and in all but one using the waveform range method (Table 4.4). Note that this was a comparison between the time points within the two template matching methods, rather than relative to the well isolated single units from principal component clustering that were used to form the templates, and so the increase was not simply a result of the method used.

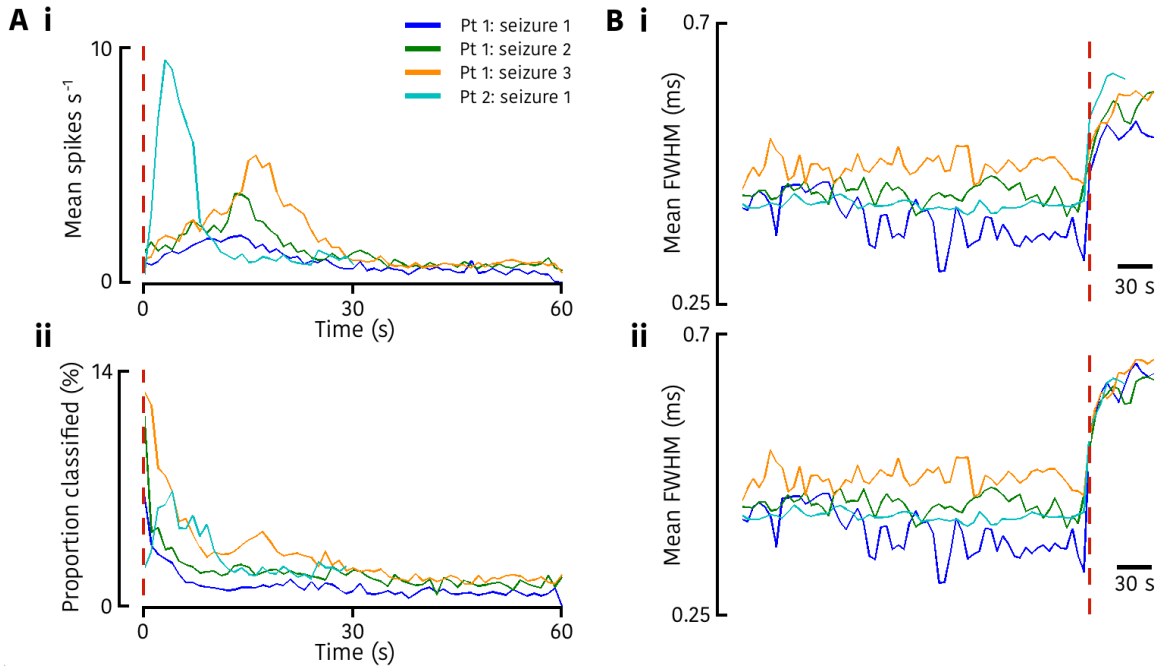


**Figure 4.7. Template matched unit features through time.**
**A.** (**i**) The mean raw number of spikes classified as being within the feature space of the well defined pre-ictal units during the ictal epoch in seizures 1 to 3 in patient 1 (blue, green and orange respectively), and seizure 1 in patient 2 (cyan; colours maintained throughout). The number of spikes was calculated in 1 second bins. Note the bias towards the early phase of each seizure. (**ii**) The mean number of spikes as in A (i), though normalized by the total number of detected spikes during that bin. In normalizing by the total detected spikes, it can be seen that after the initial seconds after recruitment (red dashed line), the number of spikes being classified as belonging to pre-ictal unit drops quickly to a simple proportion of the number of spikes occurring in the signal. **B.** The mean spike full width at half-maximum (FWHM) from the template matched units using (**i**) the waveform range method, and (**ii**) the convex hull methods. These were calculated through time from 5 minutes pre-ictal to 1 minute after recruitment to the ictal core, with the 5 minute pre-ictal results also calculated on the template matched waveforms, and so alterations during the ictal period are not a result of the change in method of spike sorting. Note the sudden increase in FWHM at recruitment to the ictal core (red dashed line); a feature that is maintained across all seizures from both patients. Spike FWHM were calculated in 5 second bins.

| Patient | Seizure | Refractory Period Violations | | | | | |
|---|---|---|---|---|---|---|---|
| | | *Waveform Range* | | | *Convex Hull* | | |
| | | Pre-ictal | Ictal | Change | Pre-ictal | Ictal | Change |
| Patient 1 | Seizure 1 | 0.23 ± 0.53% | 0.46 ± 1.31% | 204% ($p$ = 0.086) | 0.24 ± 0.54% | 1.73 ± 1.86% | 718% ($p$ << 0.001) |
| | Seizure 2 | 0.11 ± 0.30% | 0.60 ± 1.08% | 528% ($p$ << 0.001) | 0.12 ± 0.30% | 2.31 ± 2.36% | 1952% ($p$ << 0.001) |
| | Seizure 3 | 0.22 ± 0.43% | 0.82 ± 1.17% | 377% ($p$ << 0.001) | 0.22 ± 0.43% | 2.83 ± 2.24% | 1309% ($p$ << 0.001) |
| Patient 2 | Seizure 1 | 0.14 ± 0.39% | 1.38 ± 2.49% | 967% ($p$ << 0.001) | 0.14 ± 0.39% | 2.52 ± 3.22% | 1762% ($p$ << 0.001) |

**Table 4.4. Increases in average refractory period violations during the ictal period in both template matching methods.**

Whilst these template matching methods have been shown to be too inaccurate therefore for accurate spike sorting through the ictal core, any waveforms that occurred during the seizure that did correspond to the equivalent pre-ictal unit, and that unit alone, would likely

be contained within the overestimation of these methods. I therefore used these populations in order to assess any alterations in the intrinsic properties of the action potentials recorded, by analysing the spike half widths and amplitudes of the waveforms from the template matched units.

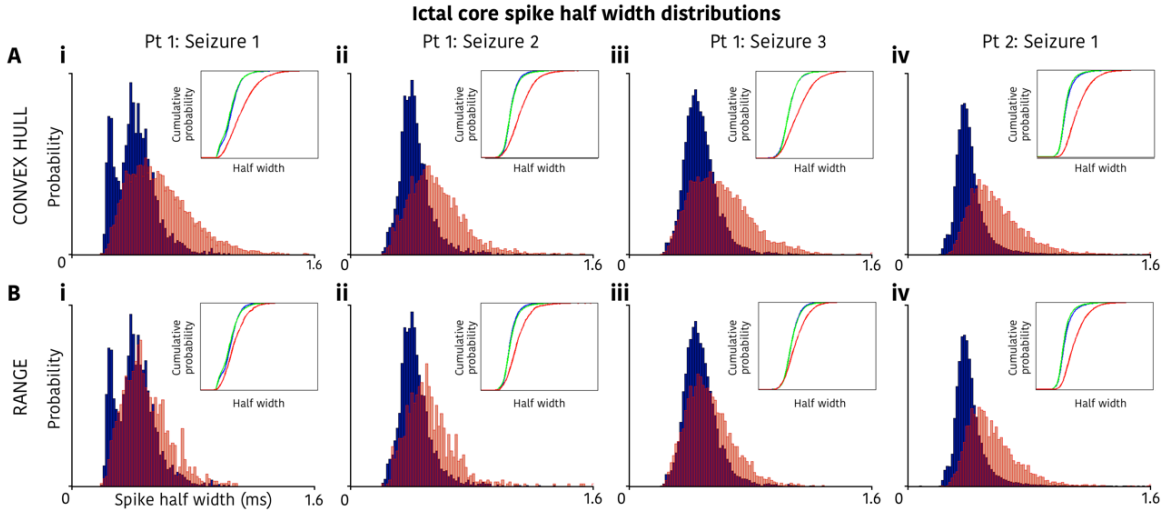


**Figure 4.8. Population increase of spike half width in the ictal core template matched units.**
Probability distributions for the spike half width of all waveforms from template matched single units, from **A.** the convex hull method, and **B.** the waveform range method, for seizures 1 to 3 in patient 1 (**i – iii**, respectively) and seizure 1 in patient 2 (**iv**). Throughout, the blue histograms show the original distribution of spike half widths from the well isolated, principal component clustered single units in the pre-ictal period. The shaded red distributions show the spike half widths from the template matched method during the ictal epoch. To show that this shift in spike half width was not a simple result of comparing the different methods, the inset cumulative probability plots show the resultant spike half width distribution of the pre-ictal epoch using the template matched method in green, along with the cumulative probabilities of the results seen in the main histograms (red and blue, colours maintained). Note that the green and blue results are very similar, with the template matching pre-ictal results (green) predominantly obscuring the original pre-ictal results (blue), throughout, whilst the ictal core template matched results are all shifted towards the right.

There was a significant increase in spike half width of units during the seizure using both methods (Figure 4.8, red histograms show the results from the seizure epoch for that method, blue histograms show the results from the principal component clustering of the pre-ictal epoch). Note that while this could be construed to be a result of the inaccurate template matching methods, the equivalent measures using the same methods in the pre-ictal epoch did not result in such an increase in half widths (Figure 4.8 insets: green shows the pre-ictal results from the template matching method; blue shows the pre-ictal results from the original principal component clustering; and red shows the template matched ictal

results), and performing the same measurements in the penumbral recordings did not show such an increase either (Figure 4.9).

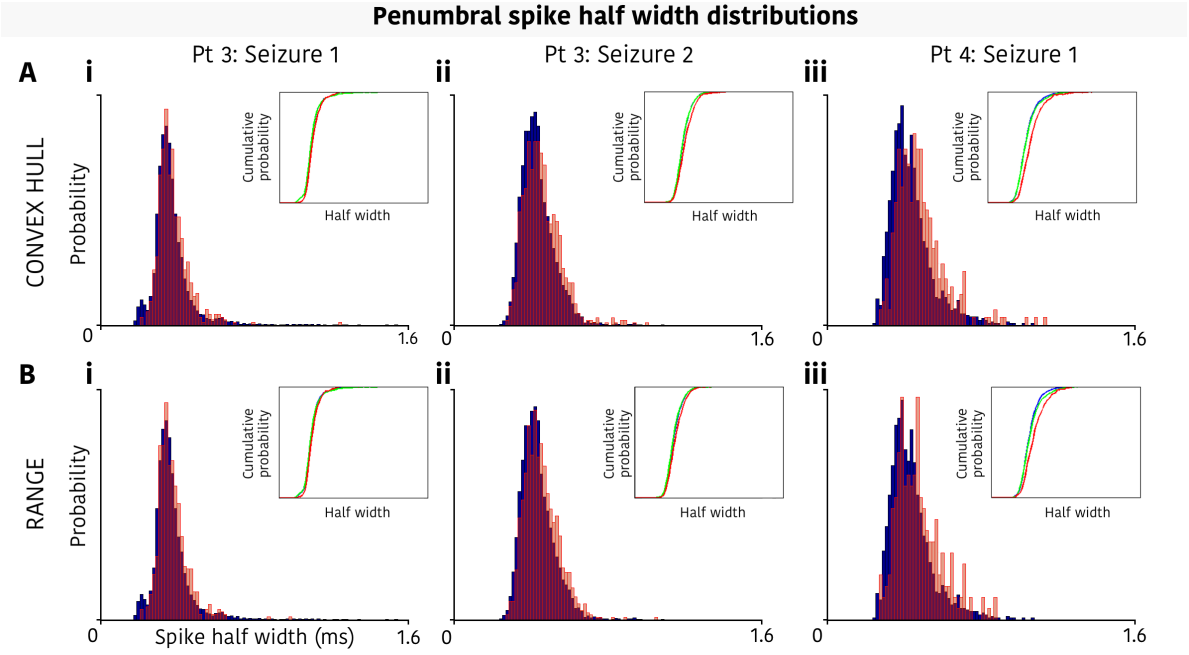


**Figure 4.9. Stability of spike half width in the template matched units from penumbral territory.**
The same format as in Figure 4.8, though with the resultant population spike half width distributions from recordings that remained penumbral throughout the seizure (patient 3, seizures 1 and 2, and patient 4, seizure 1). Note that using exactly the same methods as shown in Figure 4.8 for the ictal core, the equivalent shift in population half width cannot be seen, with both methods in the pre-ictal epoch and the template matching in the ictal epoch all resulting in very similar probability distributions, with no significant changes.

These global increases in spike half width from putative units were echoed by significant increases of spike half width within unit, with 89.3%, 84.4%, 88.2%, and 90.3% of units showing a significant ($p < 0.05$) increase (one tailed Mann-Whitney U test; mean ± SD increase of half width across all units: 133.1% ± 20.1%; 130.6% ± 17.8%; 132.4% ± 17.2%; and 136.5% ± 20.0%) respectively from patient 1, seizures 1 to 3, and patient 2, seizure 1. The equivalent penumbral numbers from patient 3, seizures 1 and 2, and patient 4, seizure 1, were substantially lower, at 11.1%, 18.9%, and 23.7% respectively (one tailed Mann-Whitney U test; mean ± SD increase of half width across all units: 98.7% ± 7.7%; 103.5% ± 9.0%; and 102.2% ± 10.6%). Furthermore, these increases in half width were found to occur immediately at seizure onset, and to remain sustained throughout, in all four seizures recorded from the ictal core (Figure 4.7 B).

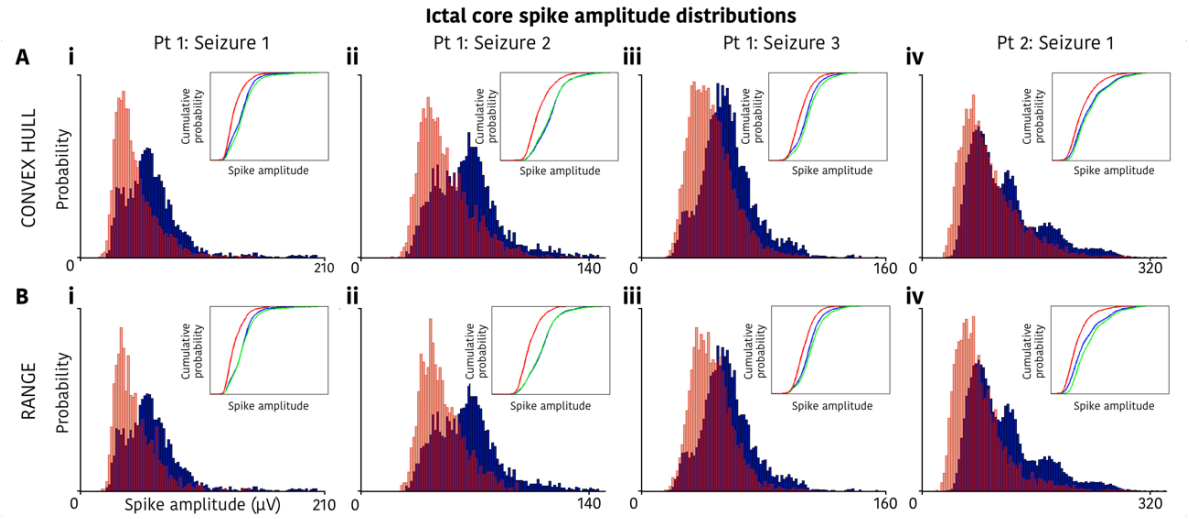


**Figure 4.10. Population decrease of spike amplitudes in the ictal core template matched units.**
The same format as in Figure 4.8, though for the spike amplitudes from all waveforms in the template matched units during the ictal epoch (red), relative to the well isolated, principal component clustered, pre-ictal units (blue). Again, this decrease in spike amplitude was not merely a result of the comparison between the two methods, as can be seen in the inset cumulative probabilities, with the template matched pre-ictal units (green) showing little deviation from the principal component clustered pre-ictal units (blue).

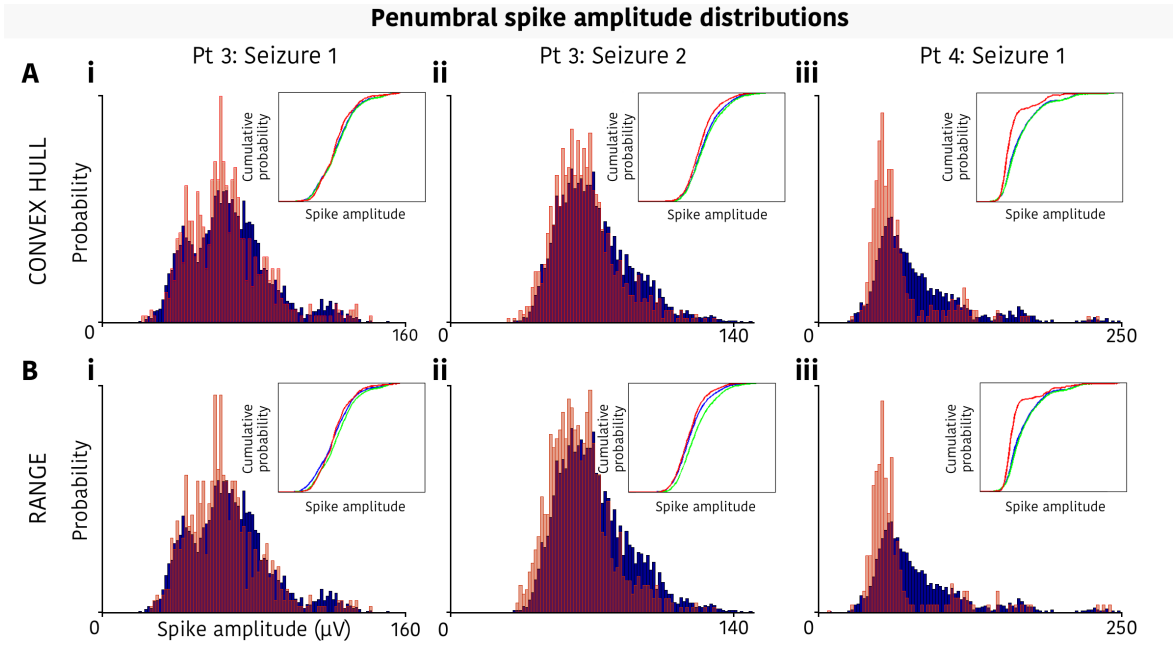


**Figure 4.11. Stability of spike amplitudes in the template matched units from penumbral territory.**
The same format as in Figure 4.8, though as for Figure 4.10, showing the resultant spike amplitudes from all waveforms in the template matched units during the ictal epoch (red), relative to the well isolated, principal component clustered, pre-ictal units (blue), from recordings that remained penumbral throughout the seizure (patient 3, seizures 1 and 2, and patient 4, seizure 1). Note the stability of spike amplitudes in comparison to the results from ictal core, regardless of method or time period. An increase in the probability of finding lower amplitude spikes during the ictal epoch was found in patient 4 using both methods (A & B iii), though the seizure in this instance lasted only 12 seconds, and so this increase in probability of lower amplitude spikes is plausibly a result of simple quiescence of the units with larger amplitude extracellular recordings.

Such a significant increase in spike half width might be expected due to nesting of hypersynchronous action potentials, in which case the spike amplitudes should concomitantly increase. Instead, I found a decrease in spike amplitude from the same template matched units in the ictal core (Figure 4.10), which again did not arise from the method alone (Figure 4.10 insets, green, pre-ictal equivalent template matching), nor in the penumbral recordings (Figure 4.11). Therefore, this broadening of spike half width is unlikely ascribable to hypersynchrony alone.

### 4.2.5. *Peri-seizure firing patterns of units in the ictal core*

Using the original spike sorting methods with the seizure epoch blanked, I next examined the recovery of single units after being incorporated into the ictal core territory in the individual patient for whom both pre- and post-ictal recordings were available (the recording in patient 2 was disrupted towards the end of the first seizure recorded, thus preventing any analysis of the post-ictal period). I performed this analysis on two equivalent duration (10 min) epochs immediately before and after the seizure, clustered independently to avoid biasing the results due to automatically matching pre- and post-ictal units as the same cell. I identified 40 cells using these shorter duration epochs, of which 30 were present both before and after the seizure, 6 cells were only present in the pre-ictal period ("lost" during the seizure), and 4 were only present in the post-ictal period. I further analysed the 30 persistent units by deriving a coefficient of drift of the centroids of the clusters in the 3D principal component space. The coefficient displayed a significant similarity when compared to the mean values obtained from comparing the relative drifts of each unit to other units at random for 10,000 trials (mean ± 2SD, observed: 0.26 ± 0.52; mean results from shuffled data (10,000 trials): 1.33 ± 2.35, $p << 0.001$, Mann-Whitney U test) (Figure 4.12 A). This shows that the waveforms found in the post-ictal period were significantly more similar in shape to their corresponding pre-ictal waveform than would be expected by chance.

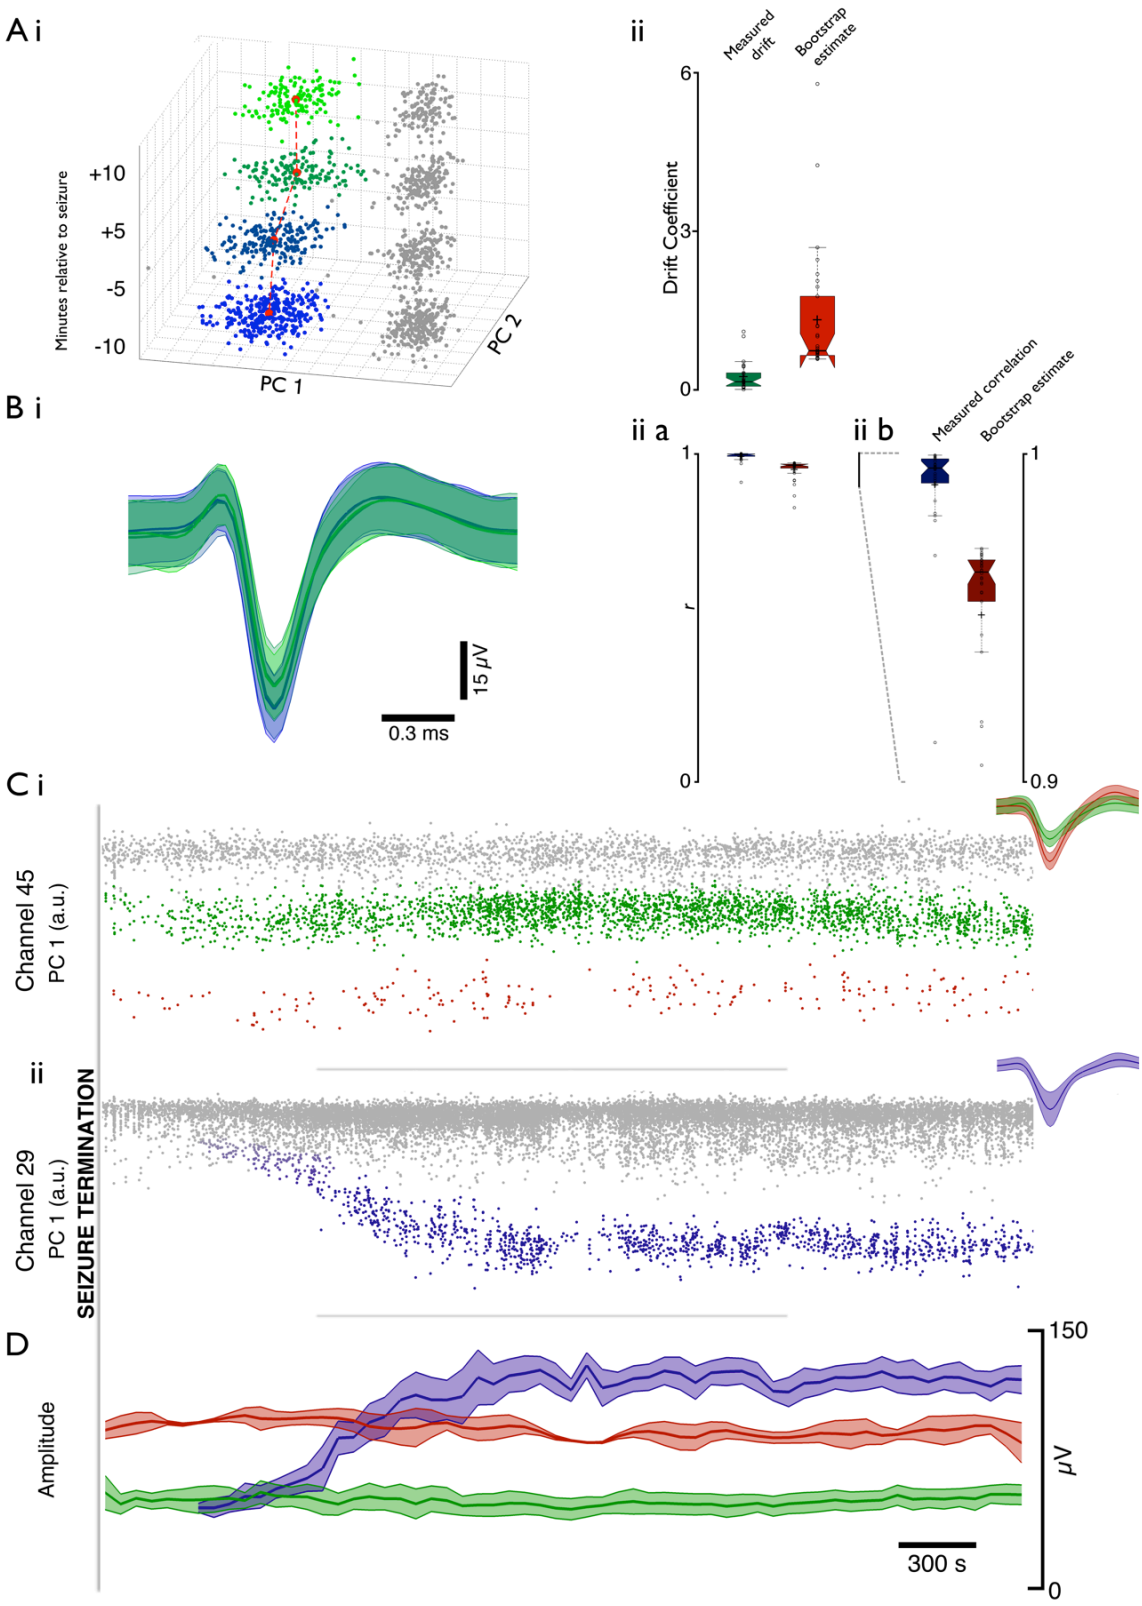


**Figure 4.12. Consistent cells evident both pre and post seizure in core recordings, with varied patterns of recovery.**

The ability to record from the same cells after a seizure as beforehand discounts movement of the MEA, due to either patient movement or vasculature response, as the cause for loss of feature specificity during seizure. **A.** (**i**) Example drift of cluster centres of a well isolated unit either side of a seizure (blue to green shows −10, −5, +5, +10 min relative to seizure onset, centre of cluster marked in red). (Cont'd on next page.)

(Figure 4.12 cont'd from previous page) (**ii**) Drift coefficient of mean post-seizure centre relative to mean pre-seizure centre (calculated in 3D space; 2D shown in i), from each cleanly separated unit with activity in both time points. Coefficient was calculated as the distance moved, divided by the prior time point's distance from zero in principal component space. *Left* (green) shows actual drift coefficient; *right* (red) shows drift coefficient as determined by comparing each unit pre to a random different unit post 10,000 times. **B.** (**i**) Mean waveforms (± 2 SD) from time points shown in A (i), colours maintained. (**ii**) Cross correlation coefficients of mean waveform prior to seizure against mean waveform post seizure. (**a**) Left, resultant correlation coefficients, and right, correlation coefficients from comparing each unit prior to a random different unit post, 10,000 times. (**b**) Spikes were detected with a negative threshold, so all mean waveforms would be anticipated to show a highly correlated shape. Expanded view of (ii a) shows significant difference between waves found pre and post relative to arbitrarily compared waves. **C.** First principal component scores over time for two example channels over a 2 hour post ictal period from Patient 1. Well isolated units are highlighted in green and red in channel 45 with unclassified units in grey. Note the immediate return of waveform after the seizure of both units. Channel 29 shows the simultaneous return of a unit with a drift in amplitude after the seizure, fading from grey to blue once distinct from the background activity. Insets show mean waveforms ± 2 SD for these clusters. **D.** Mean (± 2 SD) amplitude for the three units shown in C, over the same time period. Note the stability in the red and green units at the same time as a sizeable alteration in amplitude of the blue unit.

Further confirmation in the preservation of recordings from the same cells was found in the cross correlation of mean waveforms from each cluster, between the pre- and post-ictal epochs. Whilst a high correlation coefficient would be anticipated due to all waveforms being pre-selected based on their large amplitude negative deflections, the distribution of coefficients found was significantly different from the distribution of mean coefficients found from 10,000 trials of comparing each pre-ictal waveform at random to the post-ictal waveform of a confirmed different unit (mean ± 2SD, real: 0.99 ± 0.03; 10,000 trials: 0.95 ± 0.06, $p << 0.001$, Mann-Whitney U test; Figure 4.12 B).

The period of post-ictal quiescence varied widely between units, so to capture this, I extended these analyses up to 2 hours post seizure (seizure 1, 59 units; seizure 2, 52 units; seizure 3, 61 units). Units predominantly reappeared without any evidence for change in waveform, or position in principal component space (Figure 4.12 C (i)), however 3 of 52 units from seizure 2 and 3 of 61 units from seizure 3 (1 equivalent unit across both seizures) showed a visible slow return over many minutes of waveform amplitude and associated position in principal component space, despite the simultaneous presence of stable units (Figure 4.12 C (ii) & D).

Classifiable single units on average showed a prolonged quiescence after the final seizure discharge before the first post-ictal action potential (Figure 4.13 A; seizure 1: 123 ± 17s, 57 out of 59 units (2 did not return) from 35 channels; seizure 2: 117 ± 11s, 51 out of 52 units (1 did not return) from 33 channels; seizure 3: 157 ± 19s, 61 out of 61 units, from 41 channels), albeit with a large variance shown within this data set (seizure 1, 56-621s; seizure 2, 10-422s;

seizure 3, 15-898s). The cumulative histograms of the time to first action potentials for all units for the three seizures were quite stereotyped. Single exponential fits to these gave an average $\tau_{recovery}$ = 104 ± 22s (mean ± S.D).

Finally, I examined fluctuations in the mean firing rate across all single units before and after recruitment to the core ictal territory (Figure 4.13 B & C). There was marked fluctuation with time in the half hour before the seizure, and from this I estimated the variance of activity in this population of identified single units, with reference to the mean firing rate over that entire half hour period. Following each seizure, there was a gradual recovery of the population firing rate, returning to within 3 S.D. of the pre-ictal rate within 100s in all three seizures.

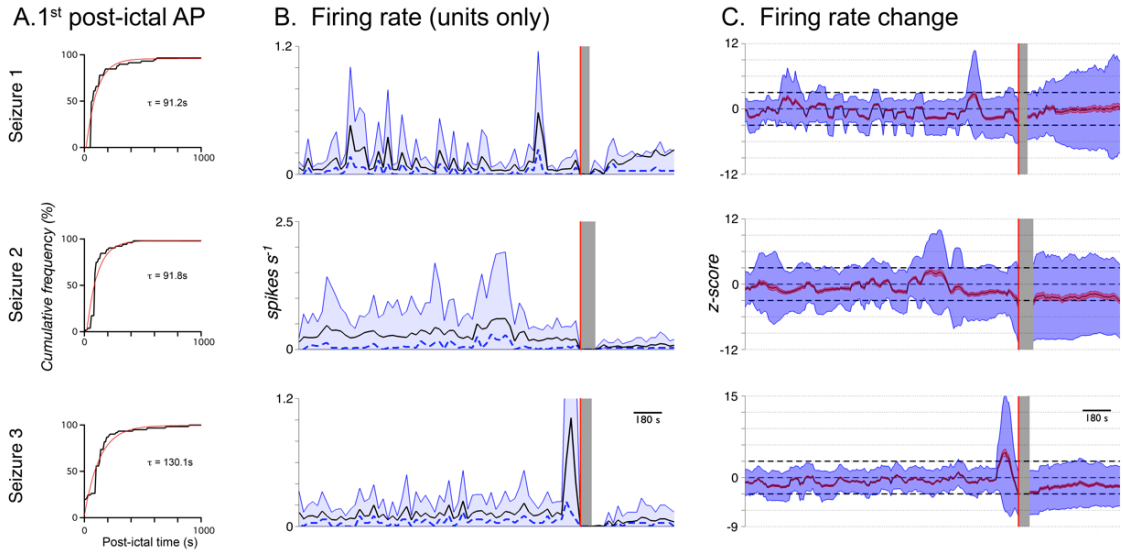


**Figure 4.13. Post ictal recovery of population firing properties.**
Firing properties for Seizures 1 to 3 in Patient 1 (recruited territory). **A.** Cumulative frequency plots for the time of first post-ictal action potential within clusters (black) and exponential fits to each (red). **B.** Mean firing rate (black), median firing rate (dashed blue) and 90th percentile of firing rate (shaded blue region). **C.** Z-scored changes in firing rate from each cell calculated in 60 s bins every 10 s, relative to the cell's firing rate during the pre-ictal 30 min. Zero therefore means no change in firing rate, with positive values being an increase, and negative values a decrease. Dashed black lines at ± 3 denote significance levels used for changes in firing rate. Mean alteration is shown in black, mean ± SEM is shaded red, and mean ± 2 SD is shaded blue. Red line in B & C denotes seizure onset and grey area shows the seizure, which was blanked out for spike sorting purposes, used throughout.

## 4.3. Discussion

### 4.3.1. *Ictal core and penumbral territories*

In this chapter, I highlight substantial changes in recording stability during seizures in certain cortical territories (ictal core), while other cortical areas (penumbra) show stable spike shapes, as has been noted by others (Truccolo et al., 2011).

The ictal core is defined by local hypersynchronous discharges as indicated by the incidence of much larger "units" starting at the time of the ictal wavefront, and also hypersynchrony over ranges of several millimetres (Weiss et al., 2013). There is also evidence in many units of intrinsic waveform alterations, resultant of either PDSs or increased extracellular $K^+$ concentrations, as indicated by distortion of any clearly separable clusters, and population alterations via template matching, although the evidence is less direct, and by analogy with comparable events in non-human recordings. The best direct evidence would come from intracellular recordings, which cannot be justified on clinical grounds.

The ictal penumbra, on the other hand, is characterised by a small increase in activity averaged across all neurons, but with a large amount of variance, with some neurons actually showing marked decreases in their firing rates (Figure 4.14). Such activity is seen immediately before the invasion of the ictal wavefront, and also in other MEAs that are never incorporated into the ictal core territories. There is no evidence of PDSs in the penumbra, and the activity there, while clearly being influenced by the core activity elsewhere, cannot be said to be hypersynchronous. Notably, this penumbral pattern of activity appears to correspond precisely with that reported in other studies of human multi- and single units during seizures, with non-hypersynchronous activity and the ability to spike sort throughout the duration of the seizure where attempted (Bower et al., 2012; Truccolo et al., 2011), and so I contend that these other studies are reporting penumbral recordings. In other words, these studies do not appear to report recordings from units that were recruited into the ictal core. Although the ictal core and penumbral territories differ in the pattern of activity, both are likely to experience compromised network function, albeit compromised in different ways (Trevelyan et al., 2013).

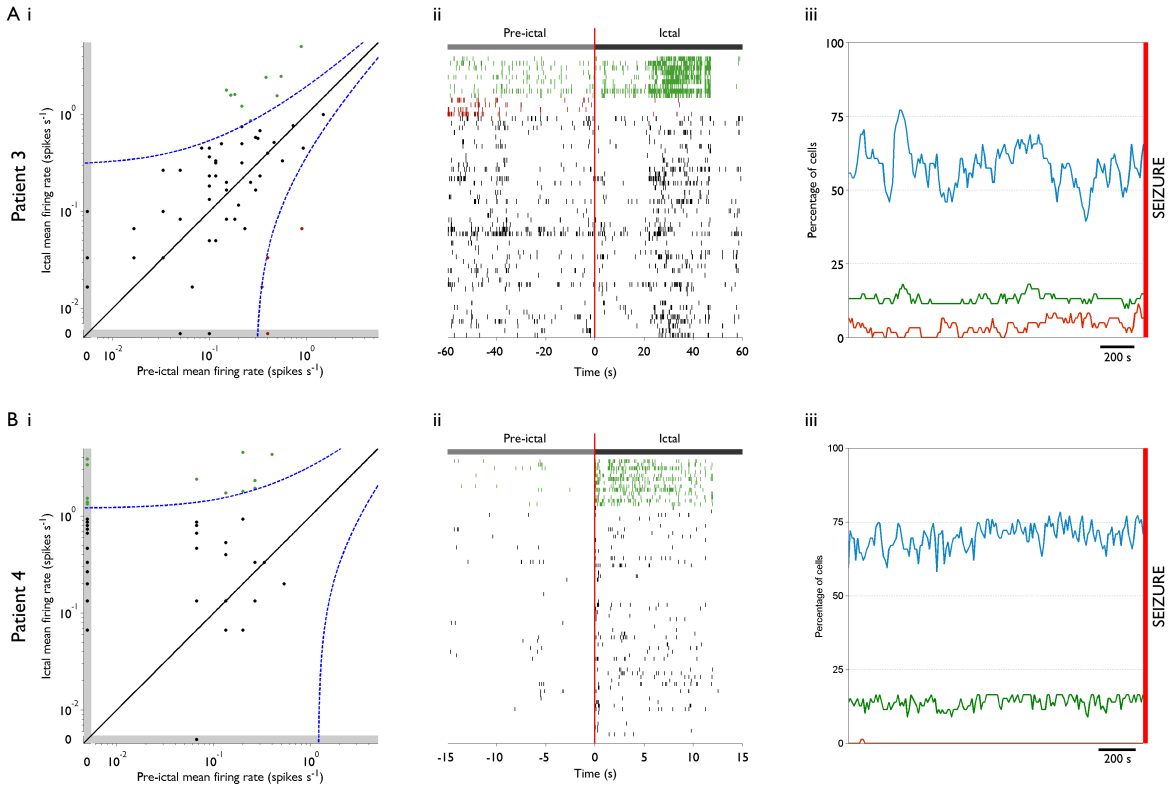


**Figure 4.14. Alteration of firing rate during seizures in the penumbra.**
**A** & **B.** Firing rates from all putative single units from patients 3 and 4 respectively (both recordings from penumbral territory), during the seizure and pre-ictal epoch of equivalent duration. (**i**) Plot of pre-ictal epoch firing rate versus ictal firing rate for all units (individual dots), on a logarithmic scale. Bounds are calculated as 3 times the square root of the firing rate divided by the duration of the epoch – an estimation of 3 SD for a Poisson distribution. Units outside these bounds were deemed to show significant increase (green) or decrease (red) during the ictal period versus the pre-ictal period. Grey bounds show units that did not fire during that epoch, and thus with a firing rate of zero, which cannot be shown on the logarithmic axes. (**ii**) Raster plot from the epochs used in (i), with units colour coded to show significant increases (green) and significant decreases (red) during the seizure, as calculated in (i). Red line shows seizure onset. (**iii**) Results from the same calculation as in (i), but taking the pre-ictal comparison epoch starting every 30 seconds throughout the 30 minutes preceding seizure onset, in order to show stability of the comparison. Blue shows what percentage of cells showed any increase at all during the seizure when compared to an equivalent epoch start at that point in time, green shows the percentage of cells that had a significant increase during the seizure compared to the epoch at that point, and red shows the percentage that showed a significant decrease during the seizure relative to that epoch.

Notably, even in the penumbra there was a period of quiescence in the single unit spikes after seizure (Figure 4.14 A (ii) & B (ii)), with complete cessation of firing across the population for 4.7 seconds other than a single spike each from 2 units in patient 3, and complete cessation of firing for 2.9 seconds in patient 4. This post-ictal quiescence has been described previously by Truccolo et al. (2011; Note that these recordings are now considered to be “penumbral” based on analyses in Schevon et al. (2012), and elsewhere in this thesis), and its presence in both the penumbral and ictal core recordings does warrant further investigation. One possibility is that both the core and penumbral territories are subject to

the same influences terminating the ictal activity. This could be a global change in metabolic state, or alternatively, it could reflect a change in drive from other parts of the brain, for example subcortical structures.

The intense shutdown of the core territories, however, represents a reduction of activity that is beyond the baseline brain state in the interictal periods, and so an alternative explanation for the quiescence in the penumbra is simply that the excitatory drive to that area is similarly reduced. Further experiments are necessary to elucidate these mechanisms, and this is a worthwhile avenue for future investigations. A first step in these procedures would be to record from two MEAs, with one in the penumbra and one in the core during seizures, if possible. These recordings would be necessary to confirm that the quiescence in both areas occurs simultaneously, though whether this cessation of firing in the penumbra is a result of lesser excitatory input after the seizure, or due to subcortical activity will require more investigation. In the meantime, we will briefly return to explore the potential role of subcortical activity in the post-ictal period in Chapter 5.

The penetrative MEAs that allow for single unit studies, sample from a small area of cortex (4 x 4 mm). In addition, the sample size of high resolution recordings is also small: we only have two patient recordings which showed full ictal core activity, although notably, all four showed periods of penumbral activity. For these reasons, the context of these recordings needs to be provided by subdural electrocorticography, which has lower temporal and spatial resolution, but more extensive coverage, and critically, is used in many neurosurgical centres worldwide. For these reasons, subdural EEG recordings probably provide the best means of testing our hypothesised distinction of seizure activity into core and penumbral territories. To this end, Weiss et al. (2013) showed previously that the onset of a phase-locked, high gamma band EEG signal corresponded closely to the time when the arrays showed a transition into the fully recruited, ictal core pattern of firing. Using these biomarkers, Weiss et al. then reanalysed contemporaneous subdural EEG recordings of old surgical resection cases to show that this new approach does indeed provide a coherent interpretation of the spread of neocortical seizures (Weiss et al., 2015). With respect to the current study, the penumbral activity recorded by the MEAs (patients 3 and 4 in this thesis) was concurrent with subdural electrode signals from elsewhere in the cortex showing evidence of ictal core activity (Schevon et al., 2012; Weiss et al., 2013). Finally, there are also clear parallels with a far more extensive animal data set. Altogether, this diverse data

set indicates a consistent association of particular characteristic features defining the ictal core and the ictal penumbral territories.

### 4.3.2. *The loss of unit specific features in the ictal core territory*

To what cause, then, can we attribute the failure of spike-sorting algorithms in the ictal core? There are a number of possibilities, which I will discuss in turn: first, that there is movement artefact, altering the relative positioning of the cell and electrode which gives rise to the unique waveform recorded; second, cell death; third, hypersynchrony results in interference between unit waveforms; and fourth, that cells experience alterations to their intrinsic wave shape, through either increased extracellular $K^+$ concentration or PDSs with a consequent reduction in amplitude of the action potential related currents. Movement artefacts are certainly a possibility during a seizure, although it was noteworthy from the simultaneous video evidence that in both the "ictal core-recordings", the loss of units occurred prior to any marked physical movements associated with the seizure.

There may also be microscopic movements of cells relative to electrodes, perhaps secondary to vascular muscle contraction or relaxation, associated with the intense neuronal activation. On the other hand, the near perfect recovery of spike shape for the great majority of units after the seizure is argument against this, because if there were significant movement, followed by recovery of the original position, this would be accompanied by local trauma, which even if small, would be expected to change the "electrical geometry" from which the unique spike shape derives. What is certain is that movement artefact is not the sole cause of the loss of units during the seizure, because in that case, the loss of particular units would be accompanied by the gain of others. Thus, there must also be contributions from the other three explanations.

Cell death cannot explain the transient loss of units, which are recovered post-ictally, but cannot be discounted as the explanation for the small numbers of unit waveforms that disappeared persistently at the time of seizures. However, features from new cells were also recorded, suggesting that array movement (minor movement, most likely), rather than cell death is the cause of these persistent changes.

Hypersynchronous bursting at the microscopic scale is readily apparent from $Ca^{2+}$ network imaging of rodent brain slices (Schevon et al., 2012; Trevelyan et al., 2006; 2007), and

sampling considerations, in which single unit recordings invariably show intense bursting activity during seizures (Ayala et al., 1973; Kandel & Spencer, 1961b; Matsumoto & Marsan, 1964), have also long argued for the same conclusion. Hypersynchrony is also apparent macroscopically in humans, since the spatial patterns of action potential firing across the MEAs show tightly synchronized discharges across several millimetres (Schevon et al., 2012). This macroscopic hypersynchrony makes it likely that at the microscopic scale too, at single electrodes, most if not all neurons are firing. The increased dispersion of data points in the principal component plots (Figure 4.3 B) can only be explained by there being constructive interference between units, creating unusually large spikes. Thus we can conclude that hypersynchrony does indeed contribute to the failure of the spike sorting.

The contribution of PDSs, on the other hand is less clear. These extracellular recordings are certainly consistent with the model proposed by Traub & Wong (1982), which indicated the prominent role of sustained voltage-gated conductances such as NMDA and voltage-gated $Ca^{2+}$ channels in shaping the PDS. These channels certainly exist in human neurons, and there is evidence for the requisite synaptic bombardment to trigger plateau potentials during human seizures. The definitive proof of PDSs, however, would require intracellular recordings to be made, and this is not feasible for human recordings.

In fact, there are several reasons for expecting spike shape to alter. It may happen secondary to changes in $Na^+$ and $K^+$ distribution inside and outside cells, or by changes in cell volume or the extracellular volume. These changes may have been expected to be associated also with DC shifts in the extracellular recordings (Gumnit & Takahashi, 1965; Ikeda et al., 1996). Such a shift was only observed in a single instance here, though it is important to note that the recording set-up and high pass filter at 0.3 Hz renders the possibility of capturing such slow transitions unlikely; it is plausible that DC shifts were evident in each case, though at too slow, or too small an effect to be recorded. Patch clamp demonstrations of PDSs are also relevant to this discussion. The cell attached recordings (Figure 4.1) record ionic currents isolated from any changes in the extracellular space by the gigaseal of the patch pipette with the cell membrane. And yet these recordings still show large changes in spike amplitude. It is possible that these are caused by intracellular changes, such as changes in cell volume or intracellular $Na^+$ concentration, but whole cell recordings, which minimise the latter, also show PDSs (c.f. examples in Trevelyan et al., 2006). Other animal work shows that spike shape changes also happen during far less intense network

activity (Harris et al., 2000; Henze et al., 2000; Ranck, 1973; Steriade, Nuñez, et al., 1993). Henze et al. (2000) performed simultaneous intracellular and extracellular recordings of the same cell, to show correlated changes in action potential shape visible in both electrodes. Harris et al. (2000) analysed the reliability of spike sorting in tetrode recordings, noting that "error rates were increased by burst activity and during periods of cellular synchrony".

### 4.3.3. *Broadening and shortening of extracellular action potential traces in the ictal core territory*

The inclusion of waveforms of approximately equivalent shape during the seizure in the ictal core via template matching introduces further insight towards the presence of waveform alterations beyond simple nesting of extracellular spikes due to hypersynchrony. These methods showed a global increase in spike half width and simultaneous decrease in spike amplitude from waveforms that occupy the same principal component space or raw voltage range throughout the spike as their well defined pre-ictal counterparts. While this may have arisen by chance due to both methods of selecting waveforms being less well defined by spike shape than the spike sorting performed on the non-seizure epochs, the same methods employed in the penumbral territory during the seizure did not produce the equivalent alterations in population wave shape, with the exception of patient 4, though only in spike amplitude (Figure 4.11 A (iii) & B (iii)) and not spike half width (Figure 4.9 A (iii) & B (iii)).

It is worth noting however that the seizure epoch analysed in patient 4 was briefer than all others, at only a 12 second duration. As a result, it is possible that multiple cells simply did not fire during this epoch. The distributions seen in Figures 4.8 – 4.11 are the probability of a waveform having that amplitude when selected at random from that population, and so this seeming increase in frequency of lower amplitude spikes may arise merely from a lack of action potentials from the cells with larger extracellular spike amplitudes (which are necessarily fewer in number due to the increase in volume recorded from as the amplitude of the signal decreases), thereby increasing the probability of selecting a low amplitude spike. This situation is unlikely to account for the ictal core examples, however, as not only were these epochs longer (30 to 120 seconds range), but they also showed a large increase in detected events, not apparent in the penumbral examples.

Further to the lack of alteration in the penumbral territories, the original pre-ictal epoch on which the templates were derived also did not show alterations to the spike half width or

amplitude upon using either the waveform range or convex hull template matching method in any patient (note the similarity between the green and blue cumulative probabilities in the insets of Figures 4.8 – 4.11). Clearly, the method alone cannot explain such an alteration in intrinsic waveform properties in the ictal core found in the template matching methods.

Perhaps, therefore, this alteration may arise from the severe increase in the number of waveforms that cross threshold in the ictal core? This may be the case if the less strict criteria for matching a waveform to a single unit resulted in the inclusion of spikes that, by chance in such a large population, are of similar shape but subtly wider. However, if this were the case, one would reasonably assume that waveforms of a narrower half width or larger amplitude would be matched at a similar rate, as the template loosens the criteria in all dimensions simultaneously. As a result, one would expect the chance detections to produce a Gaussian distribution surrounding the waveform, with a larger standard deviation but a similar mean value. Instead, no change was found in the minimal values of spike half width (Figure 4.8), and a reduction in the maximal amplitudes matched was found (Figure 4.10). In both, these resulted in shifts in the mean value of these populations, rather than just spreading of the distribution about the mean.

While these alterations in spike half width, in the distribution of waveforms that most closely resembled the pre-ictal units, are indicative of features that correspond to PDSs, they cannot be conclusive based on the method, and are indirect evidence. It is possible that the much larger amplitude waveforms arising from the nested spikes of many neurons with hypersynchronous firing would result in artefactual spikes from filter “ringing”, which would plausibly result in an increased probability of detecting spikes of larger half width than from the physiological spikes from action potentials. Filter ringing, however, typically introduces only brief periods of extra spikes in the signal on the shoulder of the large event, and of quickly decreasing amplitude, and so one would not anticipate many such spikes to fall within the range necessary for the template matching – resulting in too few artefactual spikes to shift the entire distribution of half widths. Indeed, the proportion of spikes meeting the criteria during the seizure dropped over time, with the fewest matches occurring towards the end of the seizures, in the clonic phase, when the majority of large amplitude deflections and thus the highest potential amount of filter ringing would occur (Figure 4.7 A (ii)). Overall, I conclude, on balance of evidence given the good correspondence on other measures between human and animal epileptic activity, that intrinsic waveform alterations,

either from PDSs or increases in extracellular concentration of $K^+$, do indeed occur in the ictal core territory, but not in the penumbra.

This is, of course, still indirect evidence, and serves rather to increase the confidence based on the loss of distinct clusters. Further insights may be derived through the employment of a hidden Markov model method of spike sorting during the seizure. This method does not require a separate step for the detection of spikes by a voltage threshold crossing, and therefore makes no assumptions that give rise to potential false negatives. Importantly for the recordings in the ictal core, this allows for the possibility of separating nested single unit spikes into their respective components from multiple units (Herbst et al., 2008), potentially overcoming the issues of lost spike shape through hypersynchrony, though not alteration to intrinsic spike shape from either a PDS or intense increases in extracellular $K^+$ concentration. However, the hidden Markov method of spike sorting is successful based on deriving parameters for the recordings based on an initial learning algorithm. As such, the severe alteration found in the signal upon recruitment to the ictal core would result in different parameters to those from the baseline period. Therefore, parameters from the baseline would likely be inaccurate on the ictal period, and deriving parameters from ictal periods would be inaccurate due to the typically brief duration of seizures in humans on which to train the algorithm. Nevertheless, this would be a worthwhile direction for future analyses of these recordings.

### 4.3.4. *Recovery of single unit waveform characteristics post seizure termination*

Finally, I looked at the recovery of spike sorting after the seizure, and found that the spike shape and also the average firing rate across the population recovered to pre-seizure levels within a few hundred seconds. Our recovery data comes from three seizures analysed from a single patient, so should be viewed only as preliminary data, but what is clear is that the recovery of these two features is far in advance of the recovery of normal function in this subject, who experienced prolonged post-ictal lethargy with typical corresponding EEG changes (slowing and loss of faster frequencies).

The protracted post-ictal state, which typically lasts from many minutes to hours or days (Gallmetzer et al., 2004; Rolak et al., 1992; Werhahn, 2010), therefore cannot be explained by an inability of neurons within the ictal core or penumbral territories to fire. Rather, it must be explained either by protracted altered function elsewhere in the brain, or

alternatively in terms of persistently altered emergent features from the population activity in these areas. Analyses of the population temporal dynamics of the single units in the post-ictal period are therefore the topic of the following chapter.

# Chapter 5. Post-ictal alterations in single unit activity

## 5.1. Introduction

In the previous chapter we found that there were distinctive differences between the cortical territories that were recruited to the ictal core, and those that remained penumbral throughout the seizure. Even in the ictal core, however, there was a delay of only a matter of minutes before units began firing again, with many firing after only seconds, and shortly after that the firing rates of single units returned to the baseline levels from the pre-ictal half hour.

The post-ictal period, however, has long been associated with alterations in function, with post seizure psychosis and paresis both having been described in the medical literature as early as the 19$^{th}$ century (Bravais, 1827; Jackson, 1875). These post-ictal symptoms can occur over a wide range of time scales, with post-ictal psychosis typically lasting days to weeks, while others occur sub-acutely. While single unit activity during the immediate minutes to hours after seizure termination is unlikely to have bearing on the long lasting post-ictal psychosis, they likely may prove relevant to symptoms occurring in the same time scales, such as dysphasia (Privitera et al., 1991; 1996), cognitive deficits (Helmstaedter et al., 1994; Aldenkamp et al., 1996), or paresis (Rolak et al., 1992; Pearce, 1994). It has been shown that alterations during this post-ictal time period occur in between 69 and 81% of complex partial seizures (So & Blume, 2010; Kaibara & Blume, 1988; Jan et al., 2001), ranging from delirium (Cockerell et al., 1996) to hypersalivation (Janszky et al., 2007). Despite the recognition that these symptoms can play a role in localizing the seizure onset zone (Leutmezer & Baumgartner, 2002), their aetiology is still unknown.

It is plausible that recruitment to the ictal core would leave neurons with altered excitability due to possible changes in the extracellular ionic concentrations (Bikson et al., 2003; Pinto et al., 2005), extracellular pH (Lux et al., 1986), or intracellular $Ca^{2+}$ concentration activating $K^{+}$ channels (Timofeev et al., 2004). In these cases, an alteration in firing properties during the

immediate post-ictal period would be likely in the cells that had recently been recruited to the ictal core, though potentially only briefly, for example extracellular ionic concentrations return to normal rapidly (Somjen, 2004). Such early alterations may therefore be of use in localizing the seizure onset zone with high specificity, even after the seizure has finished, significantly increasing the data available on which to determine the region to resect during surgery.

Whilst the aforementioned alterations may be brief, and therefore limited both in seizure localization potential due to time constraints and in their effects on the longer lasting functional alterations, seizure activity is likely to be a powerful force for altering synaptic connectivity, and thus secondarily altering network activation patterns. Indeed, studies in both animal models and human tissue have found alterations to dendritic spine structure resultant of epilepsy (Swann et al., 2000; Wong, 2005). These alterations are typically a loss of dendritic spines in focal tissue that has been resected (Swann et al., 2000), though further alterations have been found, including alterations to dendritic spine shape and size, branching patterns, and indeed increases in dendritic spines in the focus (Blümcke et al., 1999; Campe et al., 1997; Multani et al., 1994; Belichenko et al., 1994; Isokawa, 1997). Any alterations to these dendritic properties would alter the population dynamics of neuronal networks that are within the seizure onset zone, or a region regularly recruited to the ictal core, and are unlikely a coincidental result from ongoing synaptic alterations, as in adulthood, dendritic spines are remarkably stable (Bhatt et al., 2009). In turn, these seizure-related alterations would likely result in anticipated downstream effects; for example, in the instance of hippocampal dendritic spine loss, learning and memory deficits are likely (Wong, 2005).

It is possible that these alterations at the dendritic spine level are in fact the cause of seizures, rather than a result of them (Wong & Guo, 2013). In that instance, one would not anticipate significant alterations in the post-ictal period as a result of acute changes to dendritic spines during the seizure. However, it has been shown in slice cultures of rat hippocampal cells that intense, synchronized activity results in dendritic remodelling (Casanova et al., 2013). Moreover, time lapse imaging of dendritic spines in an *in vivo*, kainate induced seizure model in mice found that seizures of long duration resulted in greater than 75% loss of dendritic spines, and even brief seizures (< 5 minutes) resulted in moderate spine loss and dendritic beading (Guo et al., 2012). As such, alterations to the

population activity in the post-ictal period, outlasting the effects of altered ionic or pH concentrations, could be expected in cortical regions that had recently been fully recruited to the ictal core.

In this chapter I therefore present preliminary analyses from the post-ictal period, relative to the pre-ictal, baseline activity, with reference to whether or not the region recorded from was recruited to the ictal core or remained penumbral. Post-ictal alterations in intrinsic firing properties for the single units were calculated, first relative to themselves, and then relative to the population as a whole. I asked if any such alterations may aid in the localization of the seizure onset zone, and in turn, if they may explain in part, the protracted post-ictal state and its varied symptoms.

## 5.2. Results

### 5.2.1. *Altered post-ictal unit correlations specific to the ictal core and not penumbral territories*

In the previous chapter, I described the return to pre-ictal levels of firing rate across the population within the first 100 seconds of seizure termination in the ictal core (Figure 4.13 C). This showed that the protracted post-ictal state could not be explained by firing rate alone in the immediate 30 minutes following seizure termination, and so I sought to characterize the temporal dynamics of the single unit firing patterns. It was found that the spike timing of single units in the post-ictal period qualitatively showed considerable alteration in their firing dynamics in the territory that had been recruited to the ictal core, with periods of bursting across the population followed by quiescence of neuronal firing; a feature that was not present in the penumbral territory (Figure 5.1). These firing patterns suggested that divergence between the ictal core and penumbral territories during the seizure persists in the post-seizure period. I therefore sought to clarify this persistent divergent state with detailed analyses of the spiking patterns.

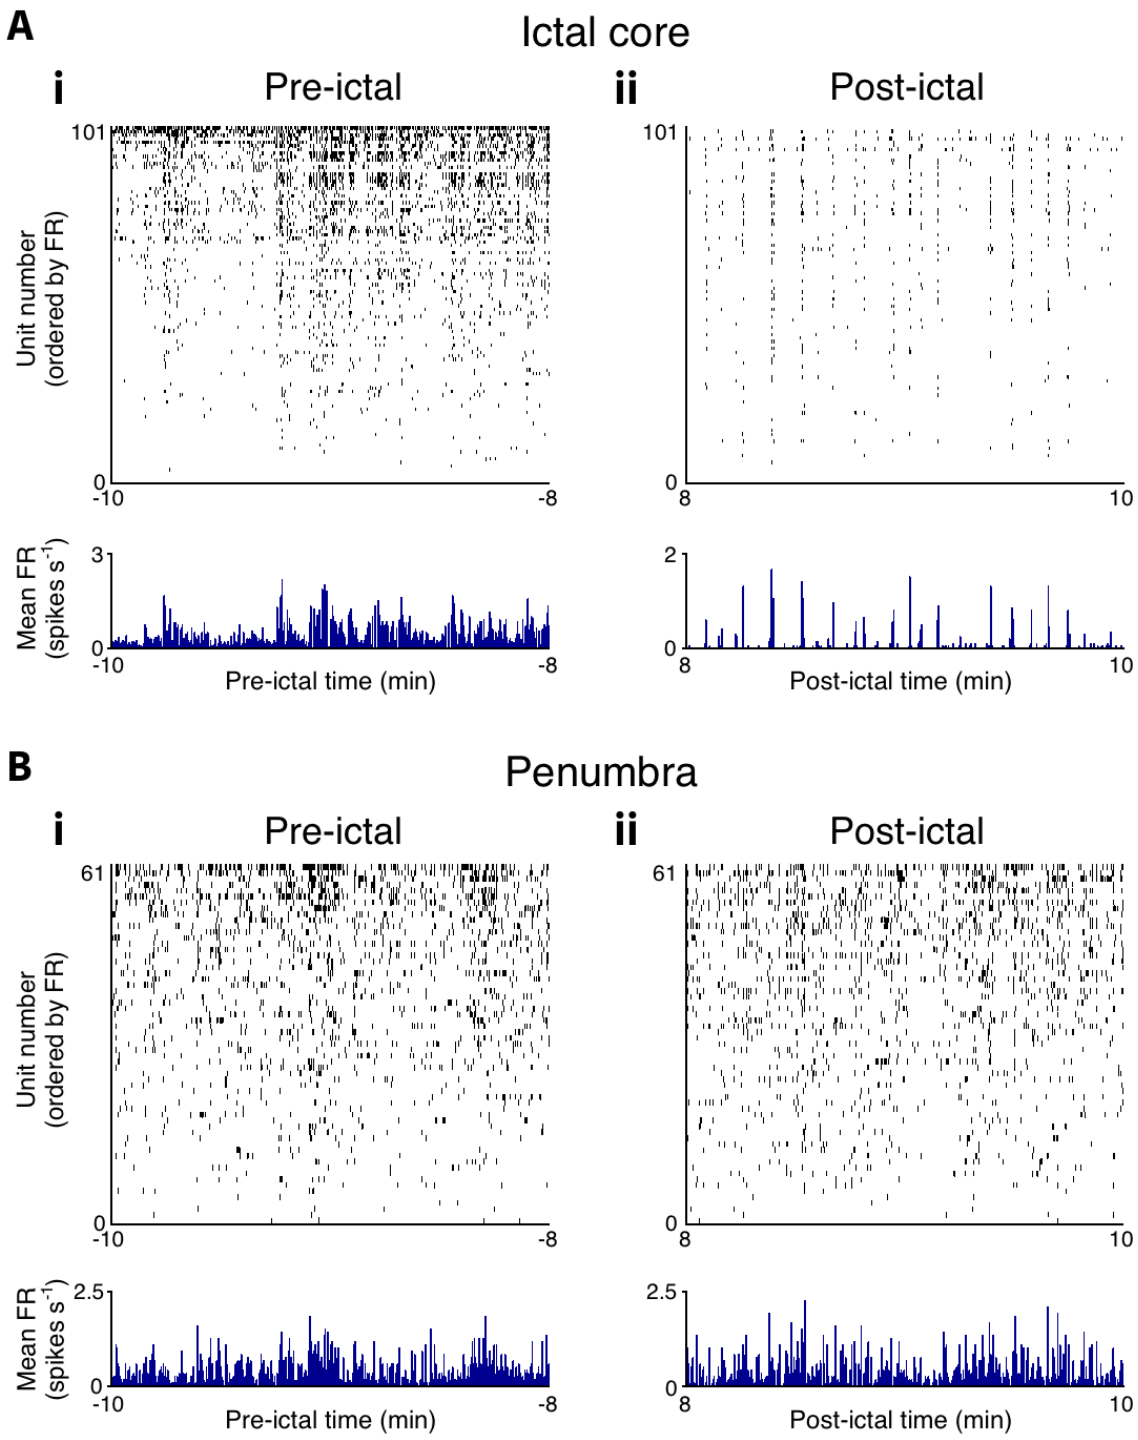


**Figure 5.1. Altered post ictal population activity in the ictal core but not in the penumbra.**
Example raster plots of firing times from all single units, ordered by firing rate, for patient 1 (**A**, ictal core) and patient 3 (**B**, penumbra), from 2 minute epochs 10 minutes prior to seizure onset (**i**), and 10 minutes post seizure termination (**ii**). Below the raster plots are the mean firing rates per unit across the population for each epoch, calculated in 200 ms bins. Note the similarity between the pre-ictal epoch in the ictal core, and both pre and post ictal epochs in the penumbra, whilst the post ictal epoch in the ictal core displays slow bursts of entrained activity across a majority of units with quiescence between bursts.

I first assessed the pairwise correlation coefficients between all single unit firing rates, as a measure of coincident firing between units. To do this, the firing rate for each unit was calculated through a convolution of the spike times with a Gaussian kernel of 1 second duration (10 ms precision, i.e. a Gaussian curve lasting 1 second, with 100 evenly spaced data points), for two 5 minute epochs, one immediately prior to seizure onset, and the other immediately post seizure termination. This provided an estimate of the instantaneous firing rate through time for each unit, upon which the zero-lag correlation coefficient was calculated between each cell pair for each epoch. An illustration of this method, along with examples of a correlated pair of units and a pair with anti-correlation can be seen in Figure 5.2. The one second duration for the Gaussian kernel with 10 ms precision allowed for a comparison of the correlation across the qualitatively seen individual bursts of activity in the population activity, whilst allowing for some variation in the exact timing of action potentials. These results were not simply an effect of the chosen kernel duration; no significant difference was found between the results from the 1 second duration kernel, and those from 500 ms ($p$ = 0.40; paired t-test), or the results from a 2 second duration kernel ($p$ = 0.27; paired t-test).

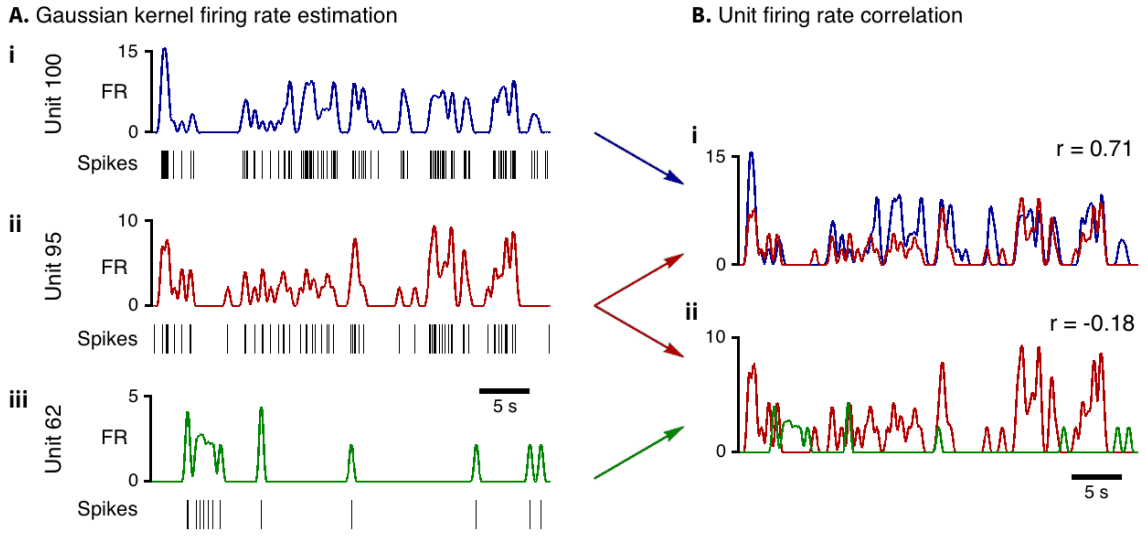


**Figure 5.2. Gaussian kernel method of estimating firing rate and example correlations between units.**
**A (i – iii).** Example raster plots of spike times for 3 example units (black, "Spikes"), with the resultant estimated instantaneous firing rate ("FR") in colour above. The Gaussian estimate of the firing rate equates to a convolution of a Gaussian curve with the pointwise spike timings. Here, the Gaussian curve has a total duration of 1 second, sampled over 100 data points, resulting in a firing rate estimation with sampling rate of 100 Hz. This conversion of pointwise spike times to a smoothed continuous signal allows for simple, direct comparisons of firing activity through the correlation between the resultant firing rates. **B.** Correlation coefficients for the example units in A, with colours maintained: (**i**) high correlation between units 100 and 95 (A (i & ii)), showing their likelihood of firing at similar times; (**ii**) a low level of anti-correlation between units 95 and 62 (A (ii & iii)), highlighting the low level of concomitant spikes between the two units.

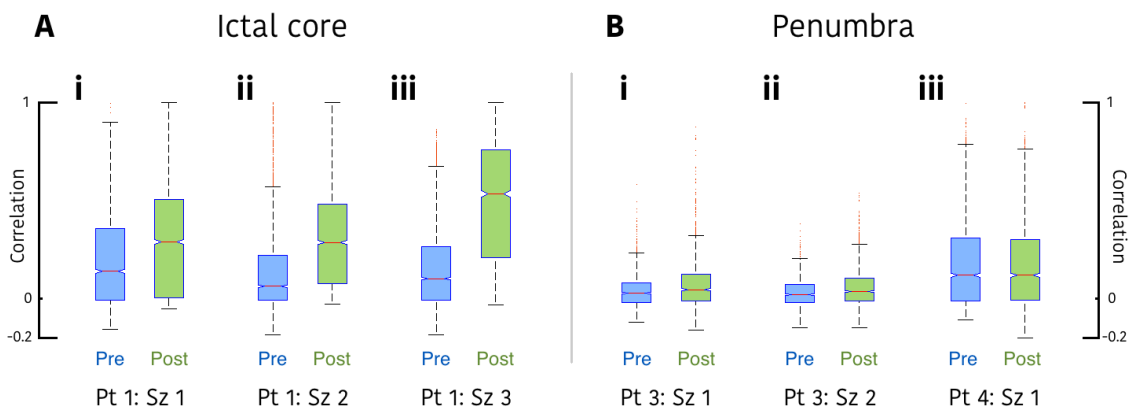


**Figure 5.3. Distribution of correlation between unit firing in the pre and post ictal epochs in the ictal core and penumbra.**
Boxplots of the resultant correlation coefficients between all unique pairs of single units from three seizures that were in the ictal core (**A**) and in three seizures in which the units remained penumbral (**B**), separated into the 5 minutes immediately pre-ictal (*blue*), and the 5 minutes immediately post ictal (*green*). The total unique pairs were 6,105; 5,050; and 4,095 for seizures 1 to 3 in patient 1 respectively (A), and 1,830; 1,830; and 3,081 for seizures 1 and 2 in patient 3, and seizure 1 in patient 4 respectively (B). These numbers correspond to every unique pairing between 111, 101, 91, 61, 61, and 79 units respectively. Note the global post ictal increase in correlation between units in the ictal core (A), which is not evident in the penumbra (B).

The post-ictal period in all 3 seizures from patient 1 (ictal core), displayed significant increases over the pre-ictal period in the range of correlation coefficients from all single unit pairs (Figure 5.3 A, Table 5.1). Furthermore, matrix plots of the correlation coefficients showing each single unit pair showed that the increased correlation was evident across a large range of firing rates. Figure 5.4 A shows the correlation coefficients for each unit pair, with the units ordered by firing rate in the post-ictal period, with higher correlations depicted as warmer colours. It can be seen that, aside from a low correlation found in the units that fired very little during the post-ictal epoch (the upper, and left bands of the right matrix plots of Figure 5.4 A), the high correlations were interspersed between units, with a slightly higher likelihood of highly correlated firing between the most active units in seizures 1 and 2. The tendency towards greater correlations in the units with higher firing rates suggests that this increased post-ictal mean correlation might purely be a result of changes to firing rate. As can be seen for seizure 3, however, beyond the quiescent units, this increased correlation was widely distributed amongst units of varying firing rates (Figure 5.4 A (iii), *right*). Importantly, also, the resultant increase in correlation between unit firing patterns could not be purely a result of firing rate, because, although there was no significant change in firing rate in the post-ictal period, a slight decrease was found in the post-ictal period for seizures 2 and 3 (Figure 4.13 C).

| | Patient | Seizure | Mean unit correlation | | | |
|---|---|---|---|---|---|---|
| | | | Total pairs (units) | Pre-ictal | Post-ictal | Mann-Whitney U |
| ICTAL CORE | Patient 1 | Seizure 1 | 6,105 (111) | 0.20 ± 0.23 | 0.30 ± 0.25 | $p << 0.001$ |
| | | Seizure 2 | 5,050 (101) | 0.14 ± 0.22 | 0.30 ± 0.24 | $p << 0.001$ |
| | | Seizure 3 | 4,095 (91) | 0.16 ± 0.20 | 0.48 ± 0.32 | $p << 0.001$ |
| PENUMBRA | Patient 3 | Seizure 1 | 1,830 (61) | 0.042 ± 0.084 | 0.074 ± 0.13 | $p < 0.001$ |
| | | Seizure 2 | 1,830 (61) | 0.029 ± 0.071 | 0.056 ± 0.099 | $p < 0.001$ |
| | Patient 4 | Seizure 1 | 3,081 (79) | 0.18 ± 0.22 | 0.17 ± 0.21 | $p = 0.86$ |

**Table 5.1. Pre- versus post-ictal pairwise single unit firing rate correlations**

To further assess the impact of firing rate on the correlation amongst unit activity, the same heat-map matrices were re-plotted with the units ordered by their firing rate during the pre-ictal epoch (Figure 5.5). Notably, whilst this did show a trend towards higher correlations between the units with higher firing rates for seizures 1 and 3 (Figure 5.5 A (i & iii)), the inverse relationship was found in the pre-ictal period for seizure 2 (Figure 5.5 A (ii)). This suggested that the relationship between firing rate and coincident firing amongst a neuronal network is dependent on secondary factors, such as brain state, rather than a simple, direct relationship. As a result, this increased correlation across units of many different firing rates in the post-ictal period, which is not present in the pre-ictal period, is likely physiological, rather than an artefact of the correlation procedure.

In patients 3 and 4, in both of whom the single units were recorded from territories that remained penumbral throughout the seizure, this magnitude of correlation was not found in the post-ictal period between pairs of differing firing rates (Figures 5.4 B & 5.5 B, Table 5.1). The overall increase was however significant in both seizures in patient 3, though no significant change was found in the post-ictal period in patient 4. The correlation coefficients from the post-ictal penumbral periods were significantly lower than those from ictal core throughout ($p << 0.001$, in each).

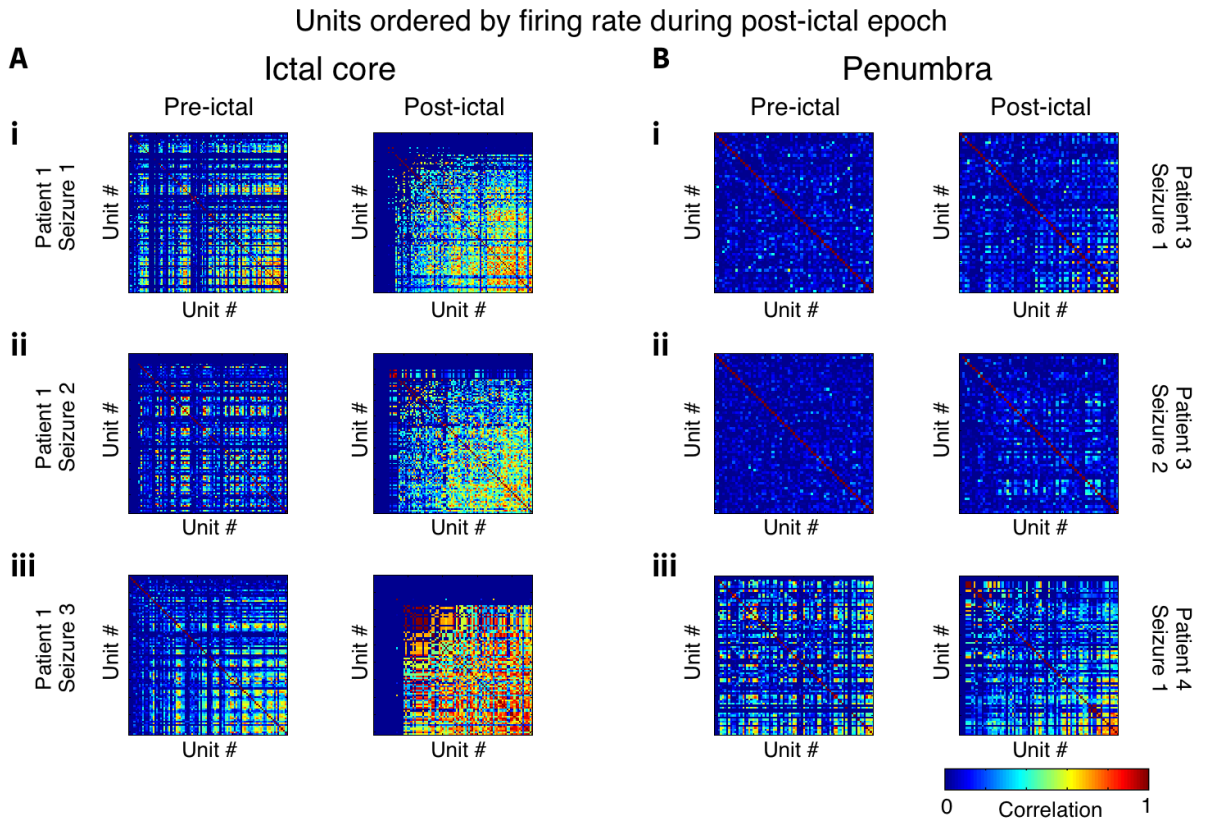


**Figure 5.4. Correlation coefficients between every unit pair in pre and post ictal epochs, ordered by firing rate during the post ictal epoch.**

Matrix plots showing the full correlation coefficient data between every unit pair that result in the boxplots in Figure 5.3, from **A (i – iii):** Patient 1 seizures 1 to 3, and **B (i – iii):** Patient 3 seizures 1 and 2, and patient 4 seizure 1. In each, the pre-ictal epoch is on the left, and post ictal epoch on the right. Unit number is ordered by the mean firing rate of each unit during the full 5 minute post ictal epoch, with the units with the highest firing rate at the bottom and right. Higher correlations are shown as warmer colours, and low correlations as cooler colours. The total units for each seizure are 111, 101, and 91 for A (i – iii), and 61, 61 and 79 for B (i – iii). Bands of low correlation can be seen in the left and upper regions of the post ictal epochs in A due to very low firing rates. Note, however, that outside these regions of quiescent units, correlation is very high between the majority of units in the post ictal epoch in the ictal core, and not in the pre-ictal epoch, or in either epoch in the penumbra. Note that while the minimal correlations in Figure 5.3 can be seen to reach -0.2, the colour range for these matrices is between 0 and 1, both for clarity, and because these data are to highlight coincident firing, and anti-correlated firing is not of interest for these measurements and can be treated as no correlation, i.e. zero in these plots.

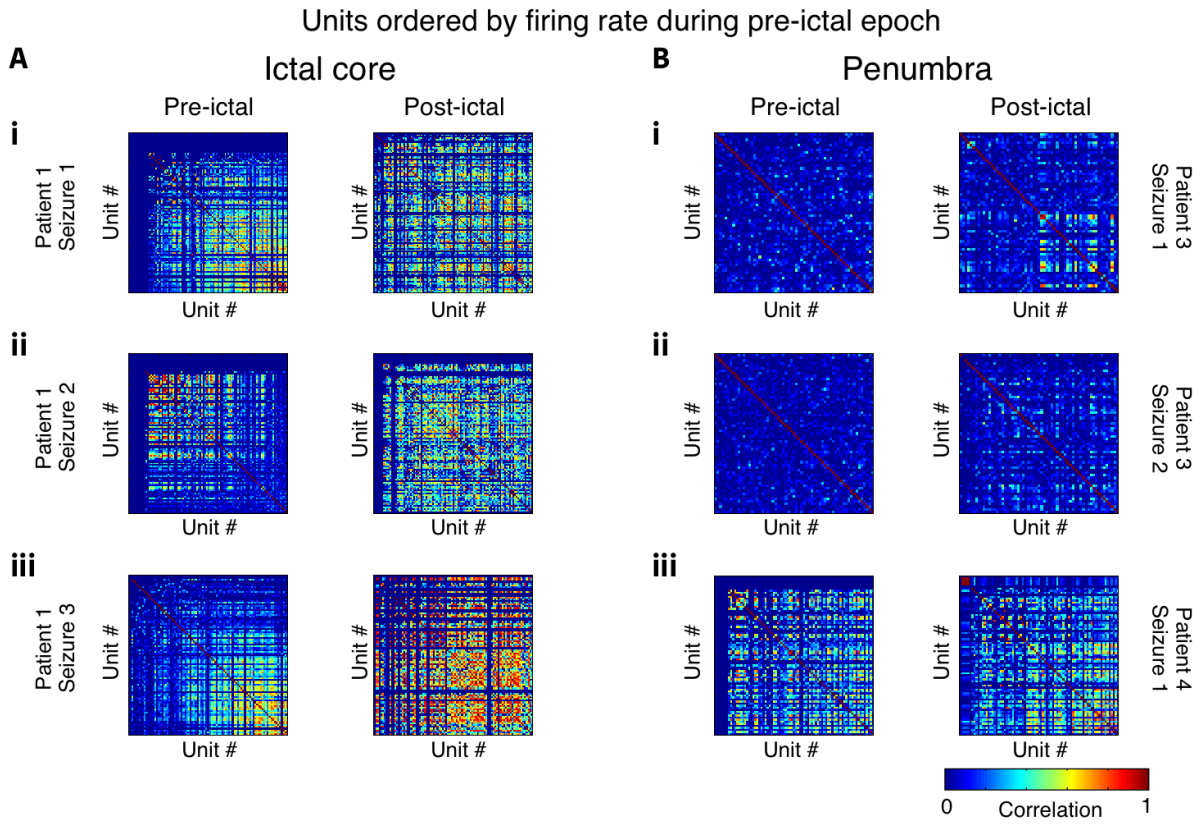


**Figure 5.5. Correlation coefficients between every unit pair in pre and post ictal epochs, ordered by firing rate during the pre-ictal epoch.**
The same format as for Figure 5.4, but with the unit number ordered by the mean firing rate of each unit during the full 5 minute pre-ictal epoch rather than the post ictal epoch. This shows that there is a relationship between firing rate and resultant correlation between units, but this is not a simple artefactual relationship. Note that in the pre-ictal epoch in patient 1, seizures 1 and 3 (A (i & iii)), the higher correlation between units is associated with those of higher firing rate, however in patient 1, seizure 2 (A (ii)), the inverse relationship is true.

I therefore looked next at alterations in the average autocorrelations for all single units between the pre-ictal and post-ictal 5 minute epochs, to assess how individual neurons' firing patterns may have altered to give rise to the increased pairwise correlations. Notably, no alterations in the mean autocorrelations were found between the pre and post-ictal periods in patient 3 for either seizure (Figure 5.6 B (i & ii)). In patient 4, however, there was an increase in the post-ictal period in spikes occurring within the 5 second window relative to that unit's other spike times (Figure 5.6 B (iii)). The pre-ictal epoch in patient 4 had noticeably fewer spikes occurring within that 5 second window than other patient's pre seizure epochs, and thus the increase may have arisen from a period of quiescence in activity in the pre-ictal 5 minutes, or alternatively, a volley of activity from many units during the post-ictal period. In fact, seizures were occurring in clusters in this patient, and this post-

ictal increase corresponds to a second seizure during the post-ictal epoch, though this does not explain the pre-ictal relative quiescence found in Figure 5.6 B (iii).

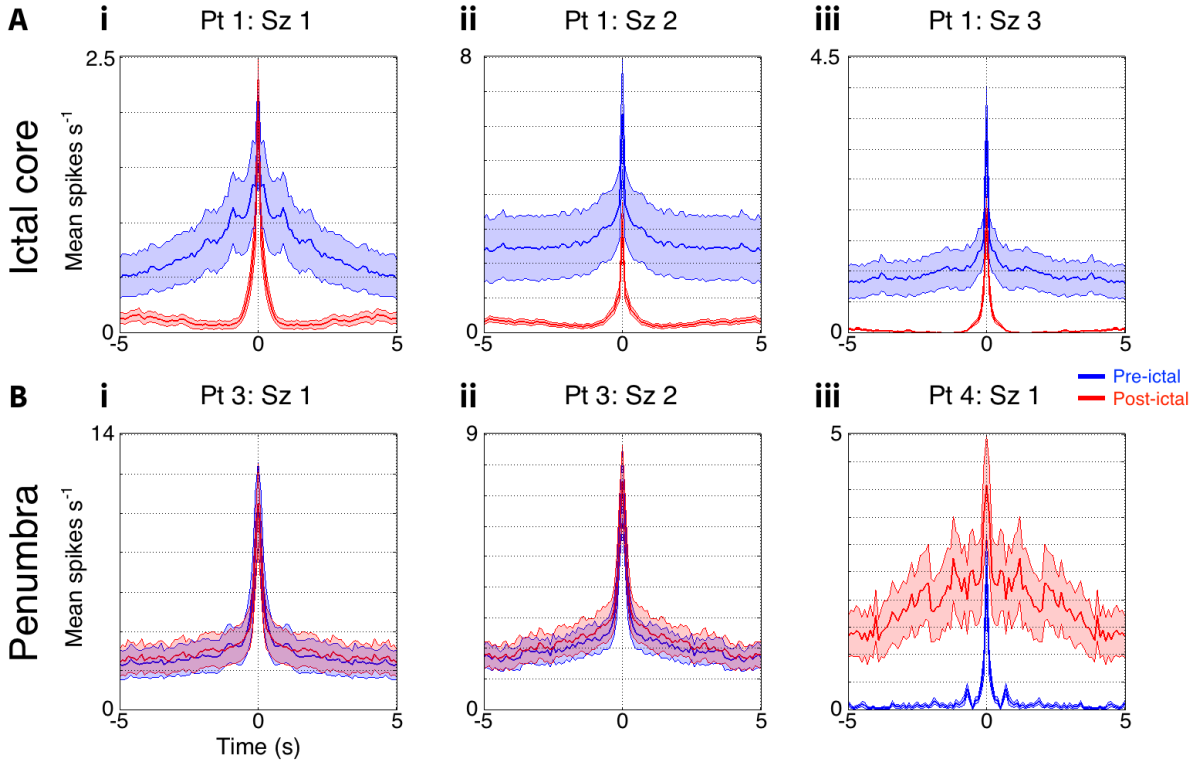


**Figure 5.6. Mean single unit autocorrelations in the pre versus post ictal epochs from the ictal core and penumbra.**

Autocorrelations were computed for each single unit over ± 5 seconds. The mean ± SD of all units from within the pre-ictal 5 minutes (*blue*) and post ictal 5 minutes (*red*) are shown for patient 1, seizures 1 to 3 (**A** (**i** – **iii**), ictal core), and patient 3, seizures 1 and 2 (**B** (**i** & **ii**), penumbra), and patient 4, seizure 1 (**B** (**iii**), penumbra). Note first the large drop in number of spikes that occur within 5 seconds of another spike in the same unit during the post ictal period in the ictal core, other than a slight probability within the immediate 500 ms (A, red); a feature that is not evident in the post ictal epoch from the penumbra (B). Note also the lack of alteration between the pre and post ictal epochs in the penumbra in patient 3's seizures (B (i & ii)), and although there is a difference between the epochs in patient 4's seizure, the quiescent epoch is in fact the pre-ictal period rather than post ictal.

In all three seizures recorded in patient 1 however, and thus in single units that were within the ictal core territory during the seizure, the post-ictal period showed a considerable drop in the mean number of spikes in the autocorrelations beyond 1 second before or after other spikes (Figure 5.6 A). Similar patterns in the mean autocorrelations were maintained for all three seizures, also with very similar patterns for all units within the epoch (note the low standard deviation in the post-ictal red shading throughout Figure 5.6 A). This alteration shows that in the post-ictal period, a putative individual neuron was less likely to fire again during a five second window after a previous action potential (or burst of action potentials occurring within a few hundred milliseconds of one another), than in the pre-ictal equivalent epoch.

Given these patterns in the autocorrelations, I next asked how this altered likelihood of action potentials in individual units in the following 5 second windows corresponded to the alterations in firing patterns of one another, as an overall change in population firing. To do this, I computed the cross correlations between the spike times of all single unit pairs, in 10 ms bins in the -5 to 5 second range, to find the mean number of spikes per second occurring in the population relative to every other spike time.

Similar results were found as in the autocorrelation calculations, with all three seizures recorded in patient 1 displaying a large decrease in the probability of spikes occurring during the ± 5 second range, though maintained probability within the ± 500 ms range (Figure 5.7 A). This showed that there were synchronized bursts of action potentials during the post-ictal period, with multiple seconds of quiescence across the population between bursts.

At the same time, very similar pre and post-ictal mean cross correlations were found in patient 3 for both seizures, indicating no major alteration in the spatiotemporal dynamics of the population in the penumbral territory (Figure 5.7 B (i & ii)). Patient 4 again showed an increase in the activity in the post-ictal period (Figure 5.7 B (iii)), a situation that was more likely given the increase found in the mean of the autocorrelations for the same units in the same epoch (Figure 5.6 B (iii)).

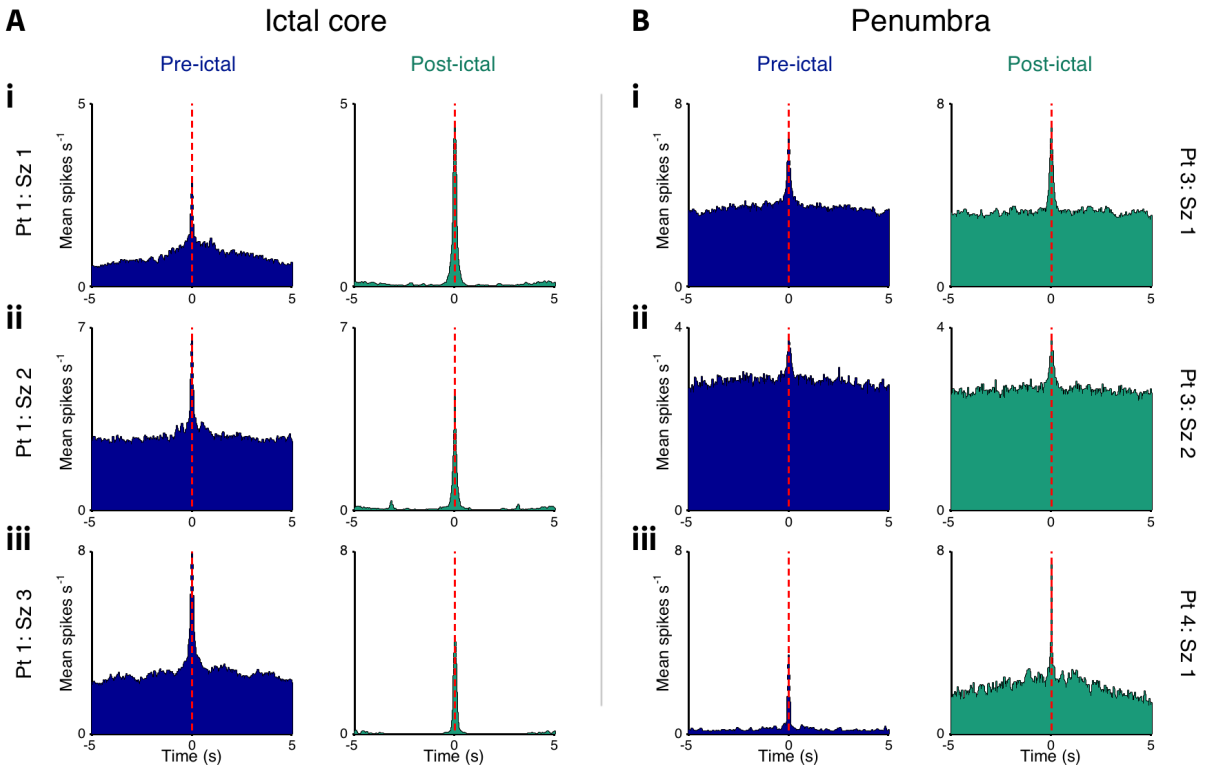


**Figure 5.7. Mean single unit cross correlations in the pre versus post ictal epochs from the ictal core and penumbra.**
Correlations between each unit pair were calculated over ± 5 seconds, as per the autocorrelations in Figure 5.6. Whilst the autocorrelations assess the level to which a unit is active relative to when it has fired before, they provide no information about how that activity corresponds to the on-going population activity. The mean results from all cross correlations for the pre-ictal 5 minutes (*blue*) and post ictal 5 minutes (*green*) are shown for patients 1, seizures 1 to 3 in A (ictal core). The equivalent results for the penumbra, in patients 3 and 4 are shown in B. These mean cross correlations show the likelihood of any unit firing relative to when a single unit fires. Axes across epochs within seizures are maintained. Note the similar features to the mean single unit autocorrelations, highlighting that these bursts of activity in the post ictal epoch from the ictal core are heavily entrained across units, with very little activity between bursts.

### 5.2.2. *Information latent in the post-ictal core and penumbral territories*

Due to the alterations in the post-ictal firing patterns in the ictal core, I next asked how this might impair cortical processing. Given that the autocorrelations and cross correlations combined gave evidence for bursts of activity occurring in the majority of units simultaneously with little activity for many seconds between bursts, it would be reasonable to assume that the amount of information that could be encoded within the population would be limited, by comparison to the less ordered structure of other time points. Therefore, for all single unit pairs, I calculated the mutual information; a measure of how much knowing one signal provides knowledge of a second. This was calculated on a full hour from each seizure where post-ictal data were available, from 30 minutes prior to seizure onset, up to 30 minutes post, data permitting. The mutual information was calculated at

zero lag between signals, using 200 ms bins to estimate instantaneous firing rate of each unit, and then calculating the mutual information on 30 second windows of data at a time, sliding 10 seconds between measurements to allow for overlap in the measurement through time. The results of the mutual information calculations were robust to bin size; Figure 5.8 B shows the results for patient 1 seizure 2, using 5 different bin widths. The only major alteration to the mutual information relative to the different bin widths was in magnitude, and when normalized to the overall mutual information across the full, 1 hour, epoch, the mutual information showed a very similar pattern through time regardless of bin size (Figure 5.8 B).

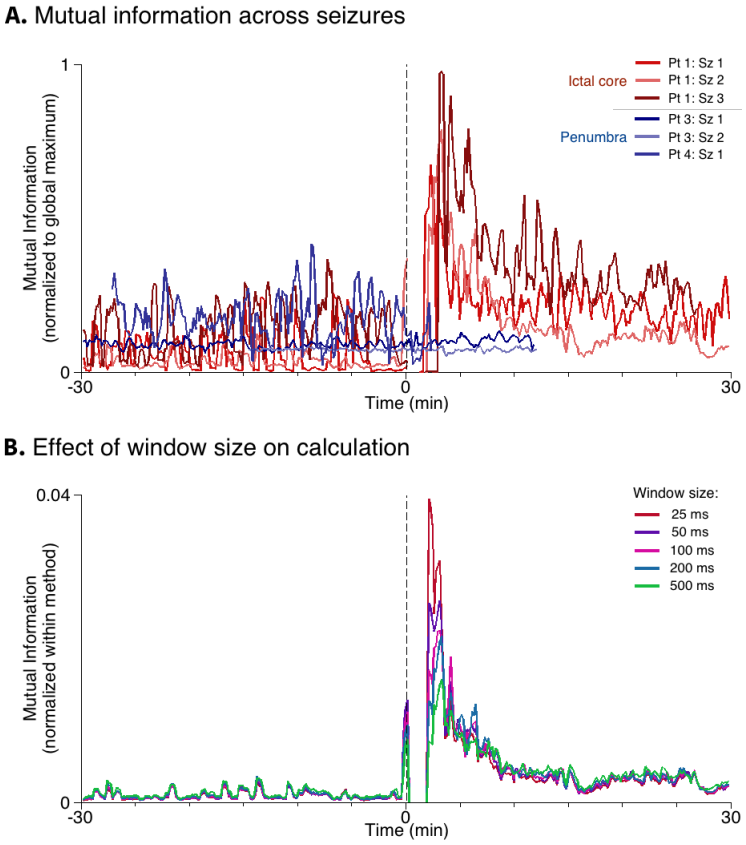


**Figure 5.8. Normalized mean mutual information differences between ictal core and penumbral territories in the post ictal period.**

**A.** Mutual information from all seizures recorded from after which post ictal data were available, normalized to the global maximal mutual information. Mutual information was calculated on firing rate in 200 ms windows. Mutual information from recordings that were within the cortical region that was fully recruited to the ictal core during the seizure are in shades of red, and those from recordings from territories that remained penumbral throughout the ictal period are in shades of blue. Note the manifold increase in mutual information in the post ictal period in the ictal core seizures (*red*), whilst the level of mutual information in all penumbral seizures is maintained with no post ictal increase. The mutual information through time is robust to the window size used, with the results from 5 different window sizes from patient 1, seizure 2 shown in **B**, including the 200 ms window used in A. Note that, unlike in A, the mutual information here has been normalized to the overall mutual information across the whole epoch within each result, rather than to the global maximum. This was done because unlike in A, magnitude is not of interest, but rather similarity of the changes through time, which are highlighted through normalization to total mutual information. Note that the mutual information is remarkably similar regardless of the window size used.

No increase in mutual information was found during the post-ictal period in recordings from units that had remained in the penumbral territories (Figure 5.8 A, *blue*). All recordings from units that had been within the ictal core during the seizure, however, displayed significant increases in mutual information in the post-ictal period (Figure 5.8, *red*). An example of the mutual information found in the population using these methods can be seen from patient 1's second seizure in Figure 5.9 A, wherein the spike times for each single unit are shown in black as a raster plot, ordered by their overall firing rates (left ordinate), with the normalized mutual information overlaid in blue (right ordinate). As can be seen, the mutual information present in the population activity shows a sharp increase immediately following seizure termination (note that the initial zero is an artefact from the blanking of the seizure epoch). Magnifications of the raster plots from all three of patient 1's seizures, from 2 minute epochs from the pre and post-ictal periods, can be seen in Figure 5.9 B (i) and (ii) respectively. The banding of action potentials in time across all units giving rise to the increase in mutual information can be seen clearly in the post-ictal period from all seizures in Figure 5.9 B (ii).

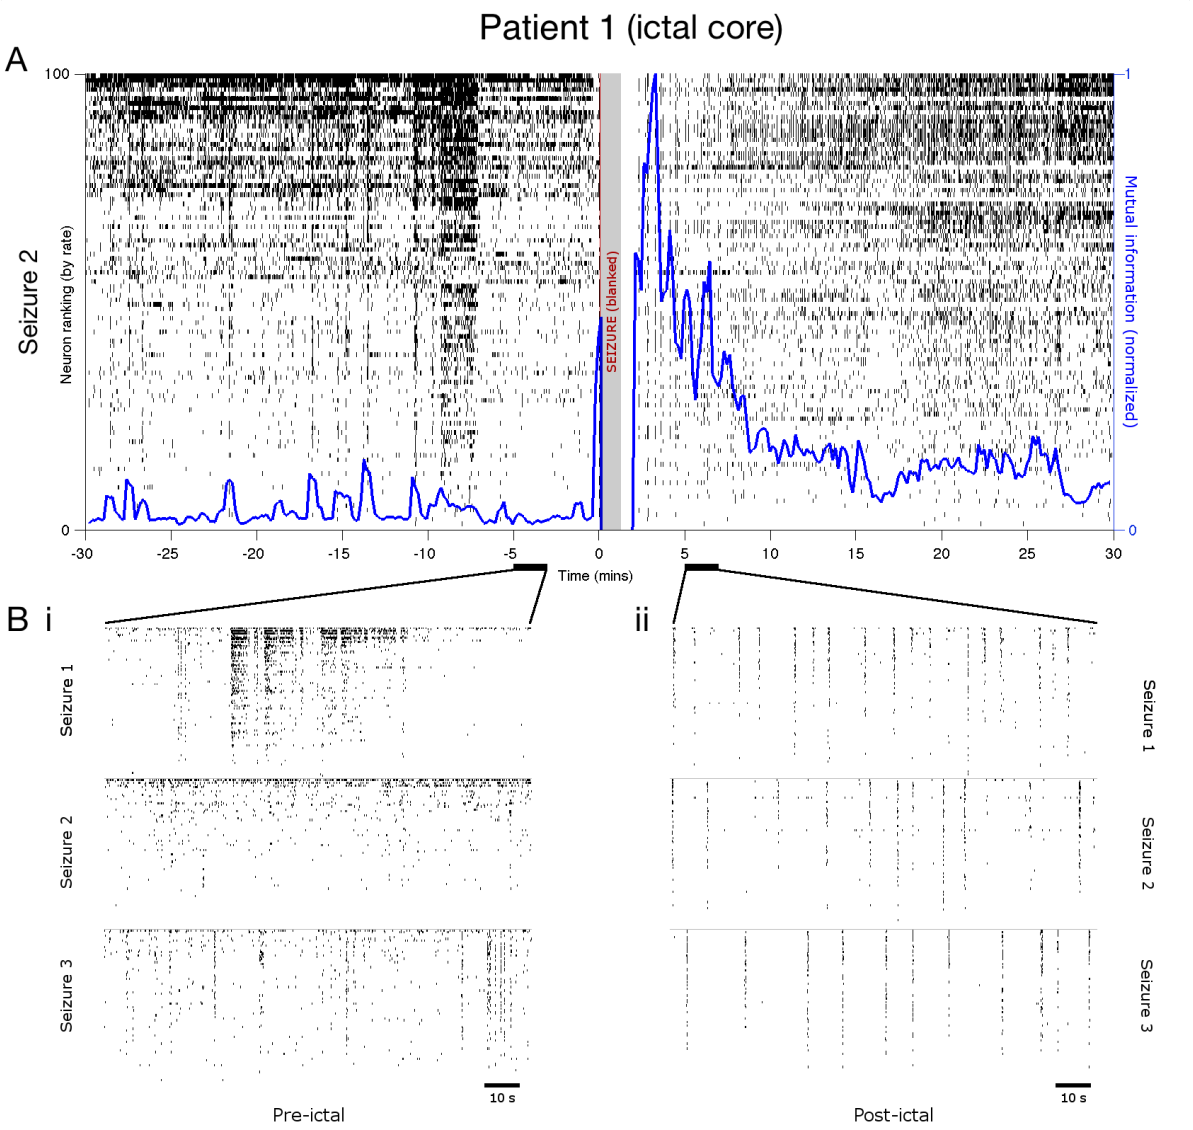


**Figure 5.9. Increases in mutual information during the post ictal period in cortex fully recruited to the seizure.** **A.** Example raster plot (*black*, left ordinate) showing spike times from all single units detected over a 1 hour period starting 30 minutes prior to seizure onset, ordered by overall firing rate. The seizure epoch has been blanked (*grey*). The normalized mean mutual information between all unit pairs from the same time point is overlaid (*blue*, right ordinate), having been calculated on 60 s windows every 10 s, on binned spike times with 200 ms resolution, as per Figure 5.8. Note the large increase in mutual information shortly after seizure termination, and the slow return to the pre-ictal baseline range. **B.** Magnifications of 2 minute epochs from a pre-ictal (**i**) and post ictal (**ii**) epoch for all three seizures recorded in this patient. Notice that the striking banding in unit firing in the post ictal period is evident across all seizures.

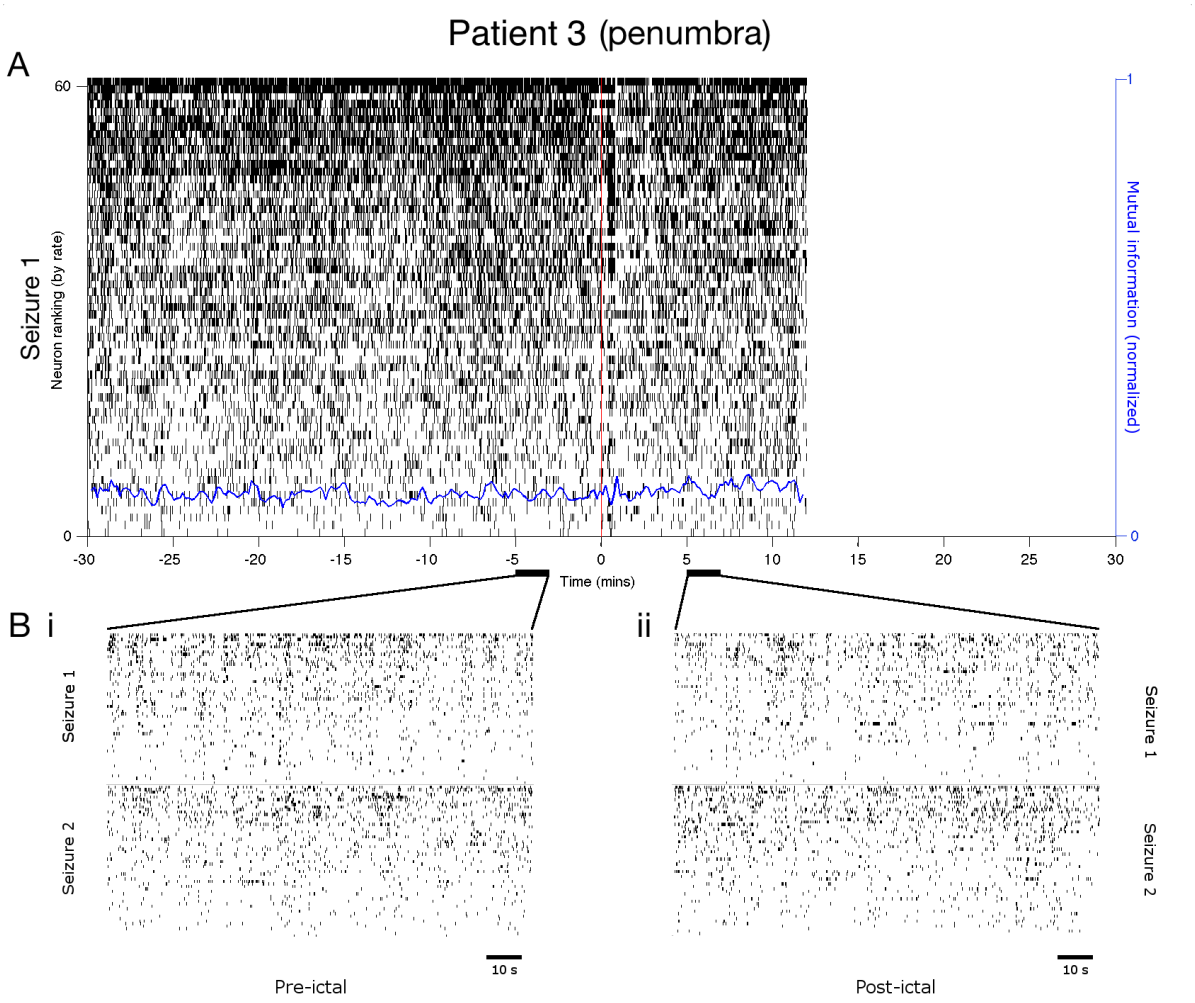


**Figure 5.10. Maintained low levels of mutual information during the post ictal period in the penumbral cortex.**

The same format as for Figure 5.9, though for patient 3, in whom the MEA remained penumbral throughout the seizure. Spike sorting was possible in the penumbra, and so the seizure epoch has not been blanked. The mutual information (*blue*) is normalized to the same scale as of that in Figure 5.9, highlighting the low mutual information throughout. Data were only available until 12 minutes post ictal in this seizure.

The equivalent measurements from patient 3 can be seen in Figure 5.10 A, displaying no such increase in mutual information, which has been normalized to the equivalent scale as for that in Figure 5.9 A. As expected, given the stability in mutual information relative to the pre-ictal period, no banding of spike times was found in the post-ictal period, unlike in patient 1 (Figure 5.10 B).

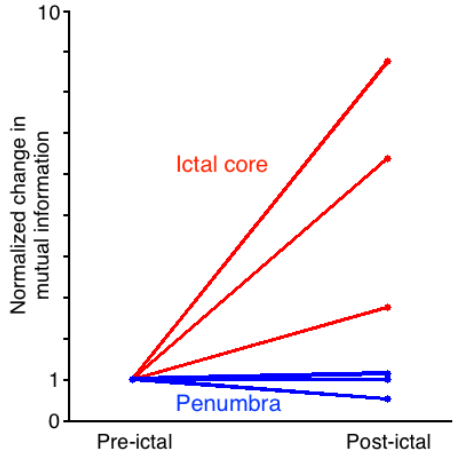


**Figure 5.11. Increase in normalized mutual information in the ictal core but not in the penumbra during the post ictal period.**
The mean mutual information from the 10 minute immediately pre-ictal and the 10 minute immediately post ictal, plotted as the change from the pre-ictal mutual information. The ictal core seizures are shown in red, and the penumbral in blue. Note that while there was a significant increase in two of the penumbral seizures, this increase was of a considerably lesser magnitude than those in the ictal core, which showed at minimum a 2-fold increase in the post ictal mutual information, and up to almost a 9-fold increase after one seizure. There was a significant decrease in post ictal mutual information in one penumbral seizure.

| | Patient | Seizure | Mutual information (bits x $10^{-5}$) Pre-ictal | Post-ictal | Mann-Whitney U |
|---|---|---|---|---|---|
| ICTAL CORE | Patient 1 | Seizure 1 | 0.52 ± 0.85 | 3.32 ± 1.83 | $p << 0.001$ |
| | | Seizure 2 | 0.46 ± 0.59 | 4.02 ± 2.09 | $p << 0.001$ |
| | | Seizure 3 | 1.83 ± 0.89 | 5.02 ± 3.28 | $p << 0.001$ |
| PENUMBRA | Patient 3 | Seizure 1 | 1.06 ± 0.12 | 1.21 ± 0.17 | $p < 0.01$ |
| | | Seizure 2 | 0.87 ± 0.059 | 0.85 ± 0.078 | $p = 0.89$ |
| | Patient 4 | Seizure 1 | 2.13 ± 1.06 | 1.10 ± 0.69 | $p = 0.99$ |

**Table 5.2. Mutual information in the pre- and post-ictal periods, with reference to pathological location**

The mutual information from the 10 minute epoch immediately post seizure termination showed a manifold increase relative to the 10 minute epoch immediately prior to seizure

onset for seizures 1 to 3 in patient 1 (Figure 5.11, *red*; Table 5.2; one-tailed Mann-Whitney U test). Seizures 1 and 2 from patient 3 also showed a significant increase in mutual information in the post-ictal epoch, though this increase was not of the same magnitude to that of the ictal core, and no significant increase was found in patient 4 (Figure 5.11, *blue*; Table 5.2; one-tailed Mann-Whitney U test). There was, in fact, a significant decrease in mutual information in the post-ictal period in patient 4 ($p < 0.001$; one-tailed Mann-Whitney U test).

### 5.2.3. *Return to baseline population activity*

I next assessed the time scale over which activity returned to the pre-ictal, baseline activity. To do this, I binned the overall firing rate from all units in 200 ms windows every 50 ms, and detected peaks in the resultant rate that surpassed the mean plus three standard deviations of the overall rate, as further described in section 2.2.3. Detected post-ictal bursts were then fitted with a Gaussian curve over a 5 second window surrounding the centre of the detected peak, an example of which can be seen in Figure 5.12. Both the intensity of the burst (amount of coincident action potentials across the population), and the related duration of the burst (spread of the activity either side of the peak in firing rate) could be quantified from the amplitude and the standard deviation of the Gaussian fit respectively. The 2.5 second duration either side of the detected peak allowed for the Gaussian curve to be fitted to the full duration of spikes that gave rise to a distinct burst of activity. At the same time, it was short enough to avoid incorporating significant "bleed through" from neighbouring bursts in the majority of cases, until they occurred close enough together in time for it to be questionable as to whether the activity corresponded to multiple bursts or rather that of continuous activity.

The bursts were found to steadily decrease in intensity, and to increase in burst width through time, with both the maximal intensity and the minimal burst standard deviation occurring within the first 5 minutes after seizure termination (Figure 5.13; the heat map colours were derived as depicted in Figure 5.12 B & C). Using 60 second windows, sliding 10 seconds at a time, mean intensity and burst standard deviations were calculated through time (Figure 5.13 A – C (ii & iii)).

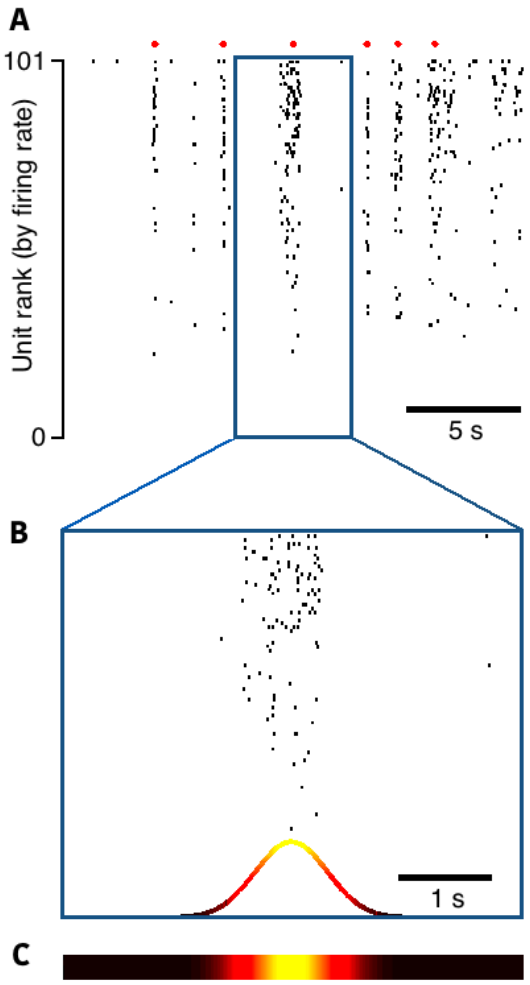


**Figure 5.12. Methods for the analysis of post ictal population bursts of activity.**
**A.** Example 20 second raster plot of bursting population activity, with units ordered by their overall firing rate. Six bursts have been automatically detected through significant increases in firing rate followed by quick return to near mean levels of firing rate (red markers above the raster plot, methods further described in section 2.2.3). **B.** Magnification of the single detected burst highlighted in deep blue box in A, over a 5 second window, 2.5 seconds either side of the centre of the detected burst. Gaussian curves were fitted to the raster plots from each burst within this 5 second window, as shown at the bottom of the raster plot. The scale for the Gaussian is arbitrary here, and chosen for clarity. **C.** Origin of the colour code found in Figure 5.13 A. The magnitude of the Gaussian curve fitted in B was colour coded from black at its minimum, and getting hotter up to white at the maximum value across all bursts. This colour coding allows for viewing alterations to the standard deviation and magnitude of these fitted Gaussian curves through time, as employed in Figure 5.13.

Although equivalent bursts were not present in the pre-ictal periods, the same burst detection analyses were performed on the 30 minute pre-ictal epoch in order to calculate the values that would be expected to occur if activity had returned to normal, and the bursting activity was no longer present. The mean and two standard deviations of these pre-ictal results are shown in green in Figure 5.13 A – C (ii & iii). Probability densities were calculated on the inter burst intervals (IBIs) and burst standard deviations from the pre-ictal, baseline epoch, from which the region outside which $p < 0.01$ was calculated (Figure 5.13 A – C (iv), yellow dotted line). The IBI range was similar between the genuine post-ictal bursts and the pre-ictal detection on regular population activity, though the burst standard deviations were considerably below the $p = 0.01$ cut off initially (Figure 5.13 A – C (iv), blue end of the colour axis).

Using the 60 second windows, sliding every 10 seconds, the delay from seizure termination until greater than or equal to 50% of bursts detected within that window were within the $p > 0.01$ range was calculated. This delay was found to be 27 minutes and 42 seconds for seizure 1 and 29 minutes for seizure 2. Due to the limited data available for seizure 3, resulting in last data available occurring at 25 minutes post seizure termination, this level was never reached for the third seizure; the maximum percentage within the cut off was 41.7% of bursts, at 24 minutes post seizure termination. Using a linear fit to the percentage of bursts occurring within the pre-ictal bounds of $p > 0.01$ through time (i.e. the number of bursts within the yellow dotted line in Figure 5.13 A – C (iv)) provided slopes of 0.69%, 0.90%, and 0.79% per minute for seizures 1 to 3 respectively. The linear fit reached 50% within their respective pre-ictal boundaries at 68 minutes 30 seconds, 50 minutes 30 seconds, and 70 minutes 20 seconds for seizures 1 to 3 respectively, with seizure 3's fit extrapolated beyond the end of data available.

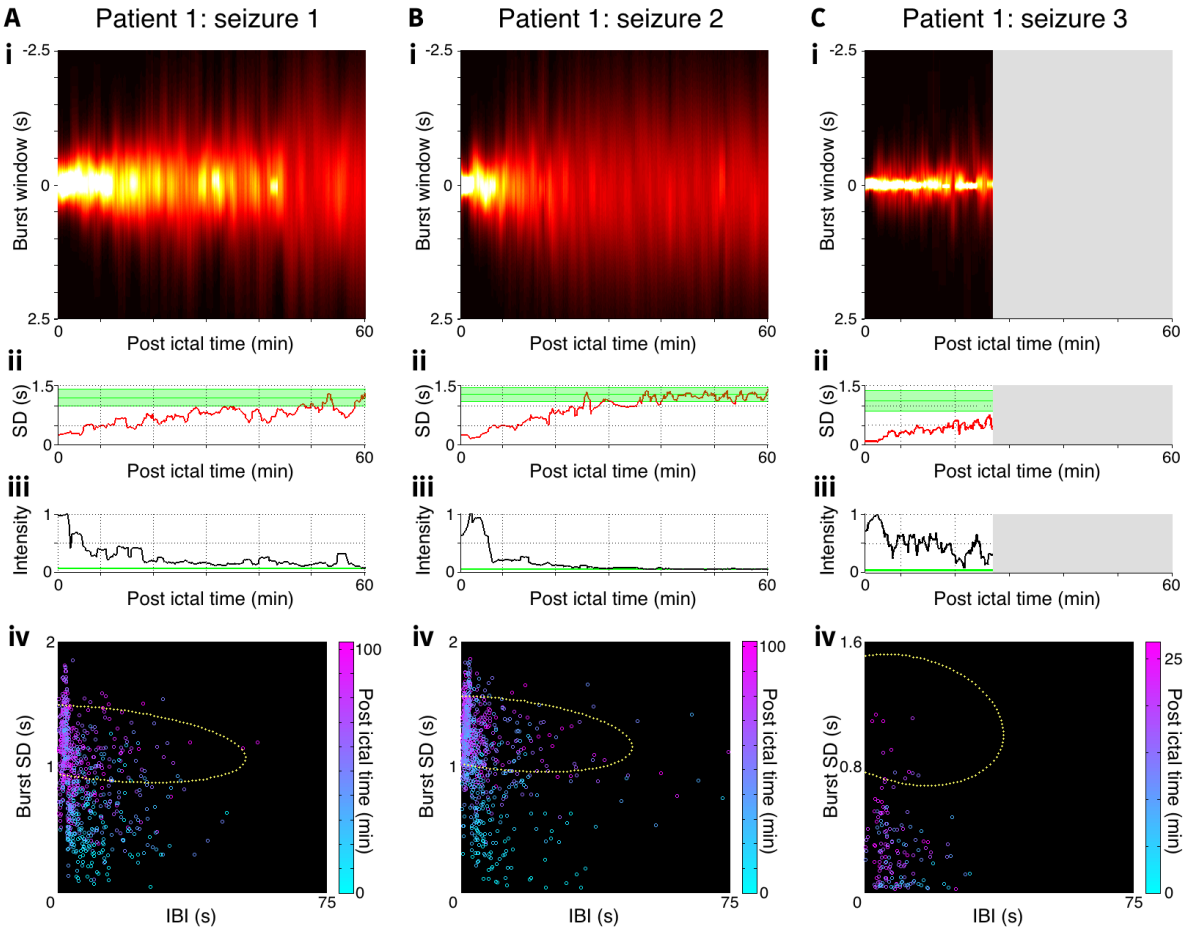


**Figure 5.13. Features of post ictal bursts through time, from all three seizures recorded in patient 1**

**A - C.** (**i**) Heat map of Gaussian fits to each detected burst of activity through time; colour axis displays the amplitude of the Gaussian curve, with the standard deviation on the y-axis, as outlined in Figure 5.12. The mean standard deviation (**ii**) and intensity (amplitude of the Gaussian curve; **(iii)**) of the bursts were calculated through time in 60 s windows every 10 s. Equivalent burst detection and associated criteria were measured on the pre-ictal epoch also, despite the lack of bursting activity, in order to measure the expected mean (green line) ± 2 SD (green shading) from a baseline, non-bursting time period. The standard deviation and inter burst intervals (IBI) for each burst were plotted through time ((**iv**), blue colours show early after seizure termination, through to purple denoting late after termination; note the different time scale in C (iv) due to limited data). Note the trend from shorter to longer burst widths through time. Probability densities were also calculated on the pre-ictal, non-bursting epoch, to detect the boundaries beyond which the probability of a burst occurring with that standard deviation and inter-burst interval fell below 0.01 (yellow dotted lines in (iv)). These probabilities were used to calculate the average delay in the return to baseline measurements of bursts.

### 5.2.4. *Local field potential correlates of single unit bursting activity in the MEA and subdural electrodes*

As these bursts were unique to the region that had previously been fully recruited to the ictal core in this data set, and showed signs of a return to baseline activity over a matter of half an hour and beyond, I next asked how the underlying lower frequency activity correlated with these bursts. I therefore close this section with some preliminary analyses of the lower frequencies in both the MEA, and the ECoG data. This was done to gain insight not only into why this activity was present in the post-ictal period, but also into potential lower frequency biomarkers that may be used to locate the equivalent region from standard subdural electrodes alone.

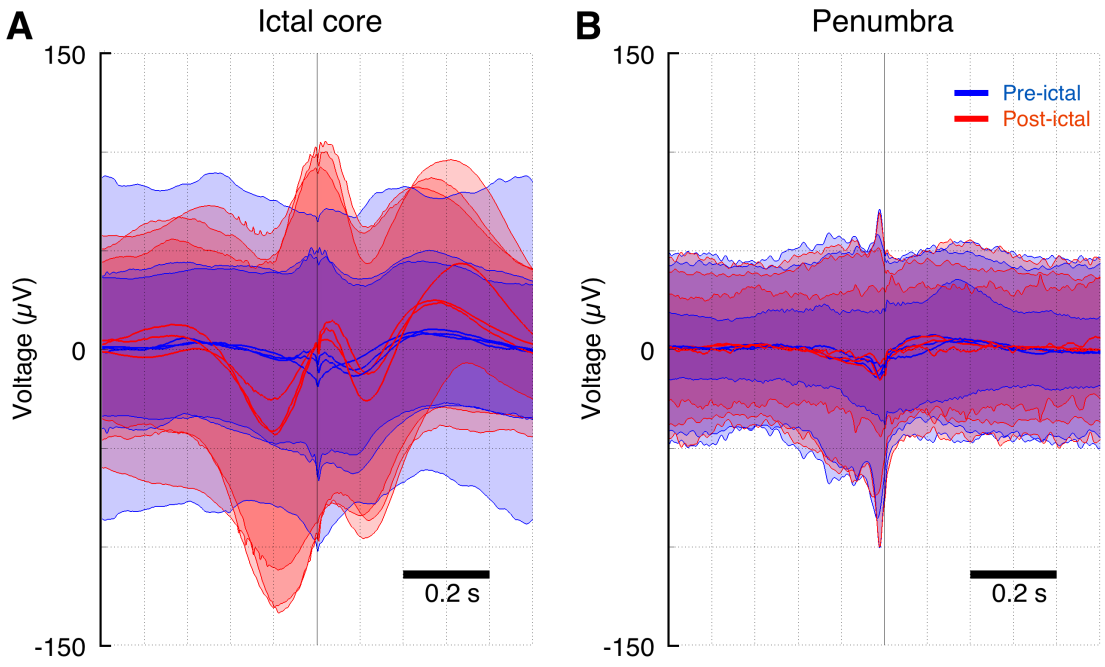


**Figure 5.14. Low frequency spike triggered averages from the MEA in the pre and post ictal periods from ictal core and penumbral territories.**
The mean ± 2 SD signals in the low frequency (1 to 100 Hz) components from the same electrode on the MEA to the spikes over ± 500 ms, calculated over all units, from ictal core seizures (**A**), and penumbral seizures (**B**). The pre-ictal spike triggered averages are in blue, and the post ictal spike triggered averages are in red. Two standard deviations are shown in the shading in the respective colours. Note the similarity between all seizures and epochs in the penumbra (B), and the similarity within epoch in the ictal core (A), whilst there is a considerable change in the spike triggered average during the post ictal period in the ictal core seizures (A, *red*), indicative of spikes riding on waves of low frequency activity.

Importantly, no qualitative features simultaneous to, and maintained across, each post-ictal burst were apparent in the lower frequencies. Therefore, the spike triggered averages, over ± 0.5 seconds, in the 1 to 100 Hz band pass filtered signal from every single unit in the immediate 10 minute pre and post-ictal epochs were calculated from the same electrode of the MEA as the unit, for each seizure in each patient, to assess if there were underlying features that were not readily apparent in single instances. Notably, all seizures from patient 1 (ictal core) showed similar pre-ictal spike triggered averages from all units in each pre-ictal epoch, and also very similar post-ictal spike triggered averages, but importantly, these two epochs' averages did not correspond to one another (Figure 5.14 A). Furthermore, the spike triggered averages from patients 3 and 4 (penumbral) throughout all seizures were very similar to one another, and similar between epochs (Figure 5.14 B). The post-ictal spike triggered averages from the territory that had been recruited to the seizure core uniquely displayed two troughs, one either side of the spike, with the spike occurring on the upswing between the roughly 5 Hz oscillation (Figure 5.14 A, red).

Given the distinctive properties of the spike triggered average from the MEA in this patient, I asked what the burst triggered averages from the standard subdural electrode array might look like, on electrodes of varying distance from the ictal core. The timings of the bursts as detected in section 5.2.3 were used to calculate the average activity over ± 2 seconds in multiple frequency bands ranging from ultra slow oscillations (0.1 to 1 Hz) up to high gamma (100 to 250 Hz), from multiple subdural electrodes, in patient 1, seizure 2. Seizure 2 was used first because data were available from more distal subdural electrodes, whereas only the 6 in the immediate vicinity to the MEA were available for this time period for seizures 1 and 3. While there was an increase in activity in the ultra slow, theta, alpha, beta, gamma, and high gamma oscillations in the burst triggered average, surprisingly, this was found across all electrodes analysed, including those distal from the ictal core (Figure 5.15 A).

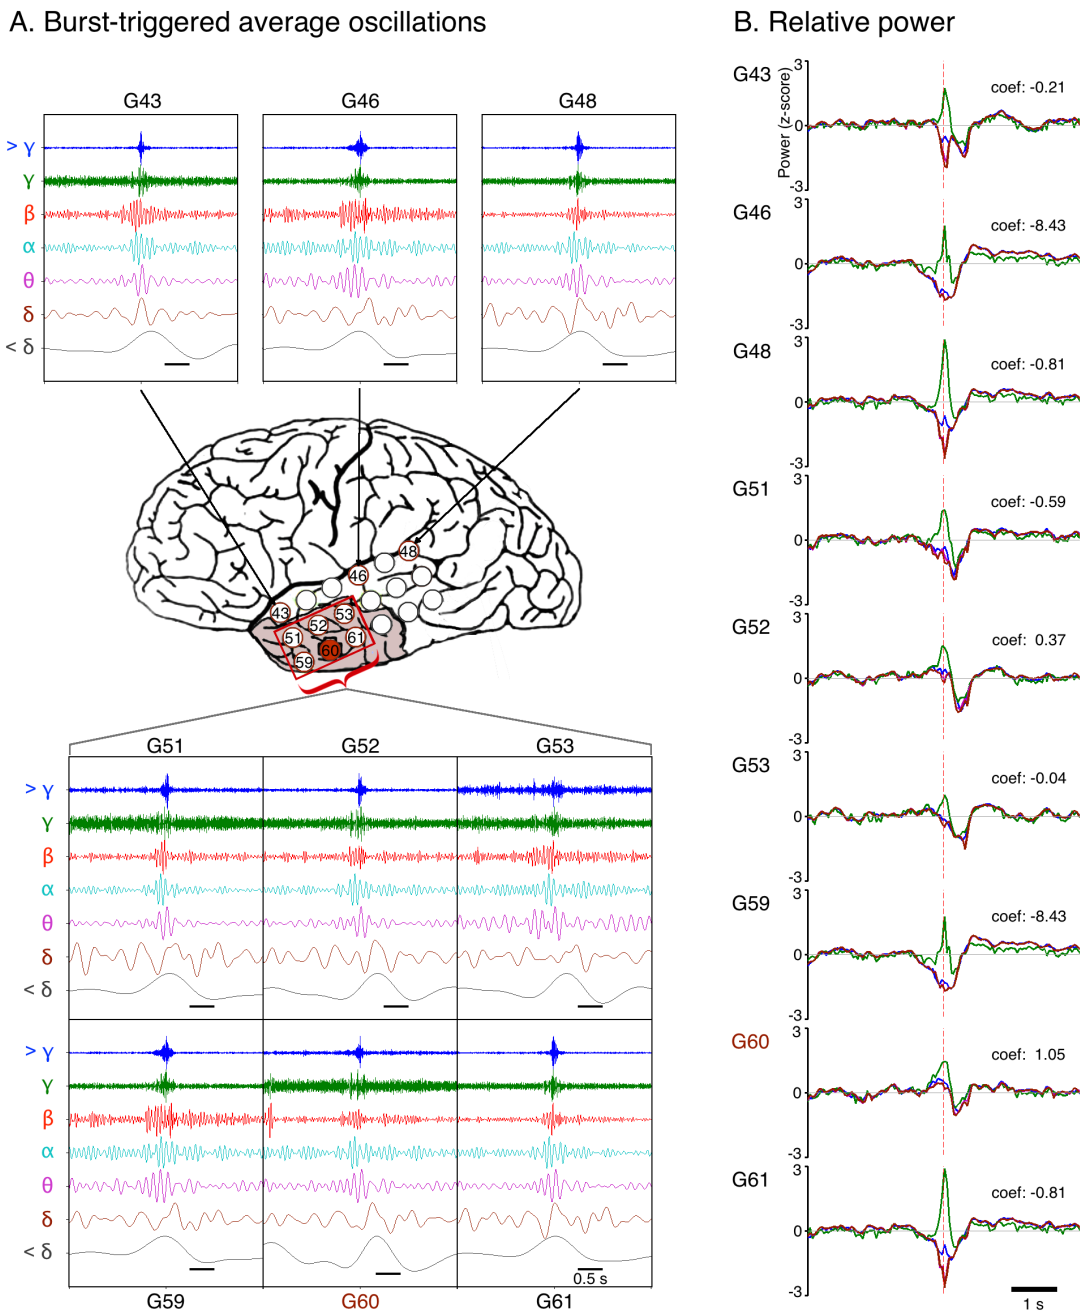


**Figure 5.15. Burst triggered oscillations in standard subdural electrodes in patient 1, seizure 2.**
**A.** Post ictal bursts of spikes in the single units were detected as in Figures 5.12 & 5.13, from which timings were used to calculate the burst triggered average in multiple frequency bands across multiple subdural electrodes. This included high gamma (> γ: 100 to 250 Hz), gamma (γ: 30 to 100 Hz), beta (β: 15 to 30 Hz), alpha (α: 8 to 15 Hz), theta (θ: 5 to 8 Hz), delta (δ: 1 to 4 Hz), and ultra slow oscillations (< δ: 0.1 to 1 Hz). The MEA from which burst timings were calculated was under electrode G60 (red). The region of proposed resection is shaded pink. Note that regardless of electrode distance from the ictal core, increases in mean oscillation amplitude were found across most frequencies. **B.** Changes in mean power (z-scored) within high and low frequency bands across bursts in the subdural electrodes. Colours are maintained from traces in A, with blue showing high gamma, green: gamma, pink: theta, & deep red: both delta and ultra slow combined in these panels. Note that in most electrodes, there is an increase in gamma power alone at the centre of the burst, with a simultaneous decrease of power in both the low frequencies and high gamma. However, in G60 (the closest electrode to the MEA) and to an extent the electrode immediately dorsal to the MEA (G52), there is a concomitant increase of power in both high and low frequencies during the burst, followed by a simultaneous decrease in power in all. A coefficient of the degree to which these powers followed this pattern was derived, shown inset. Note that the coefficients here are to be expected, as these were the training data, and are of importance instead on assessing whether the same pattern was found in the other two post ictal time points from the ictal core.

I therefore asked if there were instead distinct features of the underlying activity found across all subdural electrodes within those that had been overlying cortex that had been fully recruited to the ictal core. To do so, I assessed mean alterations in the power of the oscillations in these bands through the bursts. Noticeably, a concomitant increase in gamma oscillations was found across all electrodes at the centre of the burst, though the simultaneous power of the high gamma band, and the theta, delta and ultra slow oscillations differed by location (Figure 5.15 B). Qualitatively, the subdural electrode closest to the MEA, and the ictal core, appeared to show a coincident increase in power across all bands mentioned, followed shortly after the burst by a coincident decrease in power in all (Figure 5.15 B, electrode G60), which was less apparent in electrodes more distal to the ictal core. I derived a coefficient to measure both the similarity of the change in power over time, and the presence of a positive mean deflection in power prior to the burst along with a following negative deflection in order to quantify whether this pattern of power alteration through time was present in the other post-ictal periods in patient 1. The coefficient was measured as the mean of the correlation coefficients between each of the frequency bands, multiplied by the positive area in the -0.3 to 0.2 second range relative to the burst centre, divided by the negative area in the 0.2 to 0.7 second range relative to the burst centre, as below:

$$coef = \bar{r} \times \frac{\sum_{-0.3}^{0.2} \bar{p}}{-\sum_{0.2}^{0.7} \bar{p}}$$

where $\bar{r}$ is the mean of the correlation coefficients, and $\bar{p}$ is the mean power through time. Note that the timings and overall design of this coefficient were created so as to fit the qualitative data seen in Figure 5.15 B, and so the resultant coefficients seen therein are to be expected.

I applied this metric post hoc to the other two seizures from patient 1, however, and found, although lowered in magnitude, the same pattern in the electrodes surrounding the MEA and ictal core. The oscillation increase was still evident across all electrodes, along with very similar power alterations through the burst in the high and low frequencies. The only positive coefficients were found in the electrode immediately adjacent to the MEA (G60), and the electrode immediately dorsal to that (G52), with all others showing negative coefficients (Figures 5.16 B & 5.17 B).

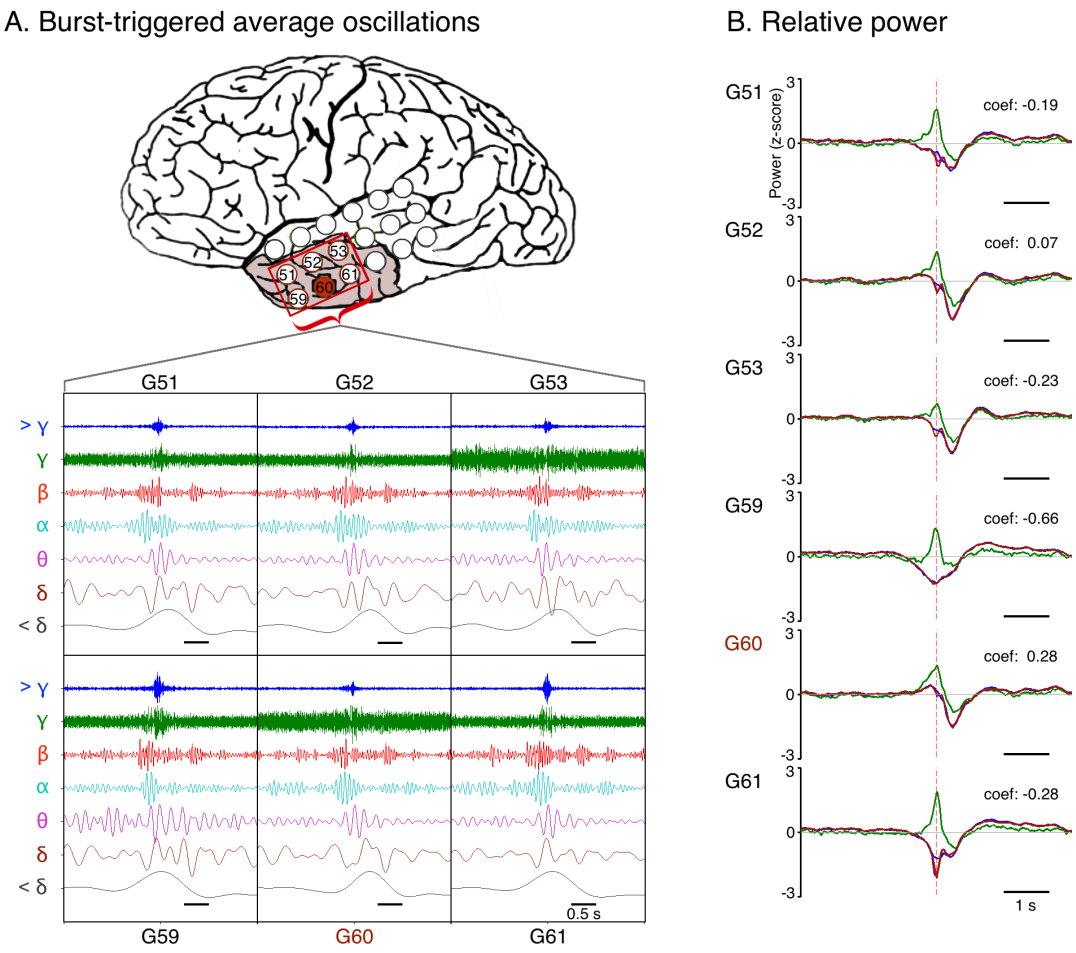


**Figure 5.16. Burst triggered oscillations in standard subdural electrodes in patient 1, seizure 1.**
The same format as for Figure 5.16, though for patient 1, seizure 1. Note the maintained features, and, although at altered magnitudes, similar pattern in coefficients as those found in Figure 5.15.

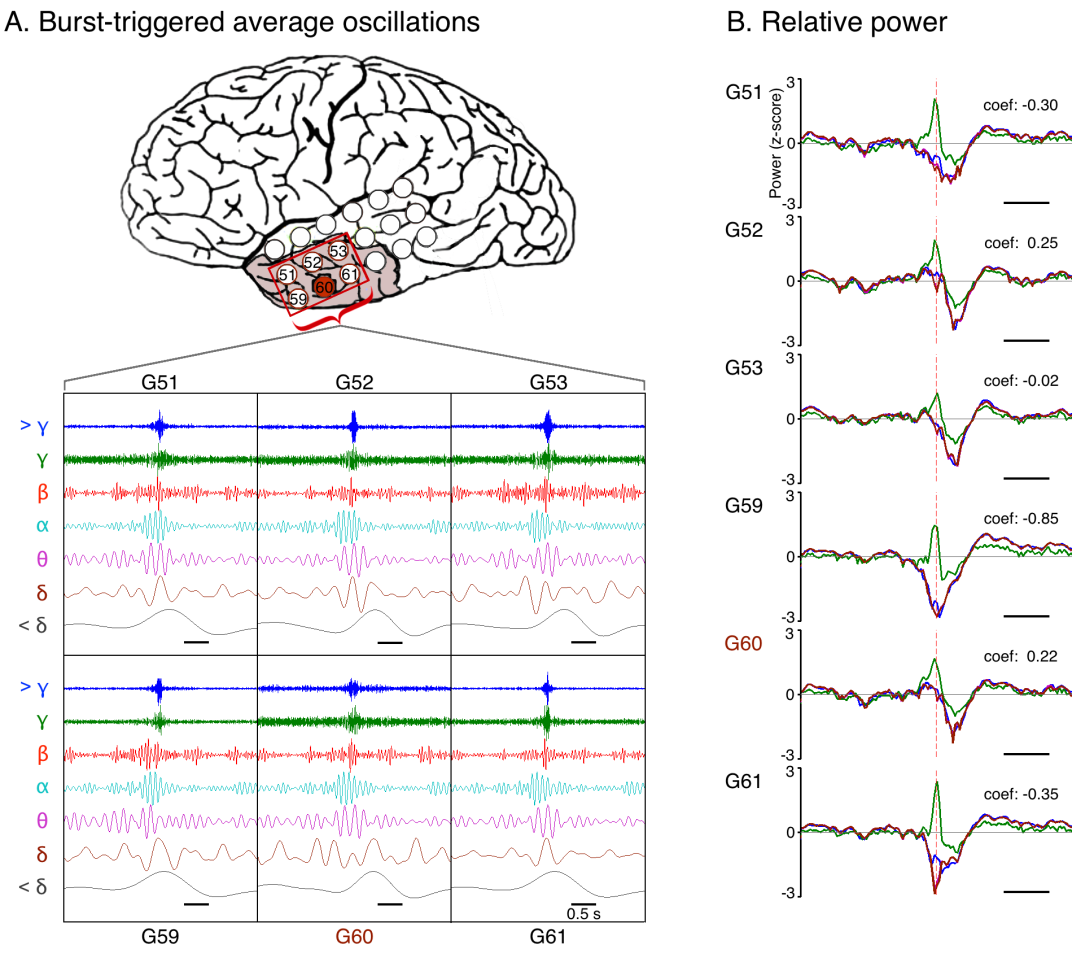


**Figure 5.17. Burst triggered oscillations in standard subdural electrodes in patient 1, seizure 3.**
The same format as for Figure 5.16, though for patient 1, seizure 3. Note the maintained features, and, although at altered magnitudes, similar pattern in coefficients as those found in Figures 5.15 and 5.16.

While enticing, it is important firstly to bear in mind that these are data from ictal recordings from a single patient as this is ongoing work, and although maintained from seizure to seizure, the underlying alterations in oscillatory power during the post-ictal bursts could arise for many reasons, as discussed below.

## 5.3. Discussion

Although not as directly incapacitating as the seizure itself, or as well studied, the time immediately following seizure termination – known as the post-ictal period – may also severely impact a patient's quality of life, especially given its long duration. In fact, the post-ictal period has been associated with a myriad of symptoms from as early as Bravais's research on post-ictal hemiparesis in 1827 (Bravais, 1827). Since then, this protracted state after seizure termination has been associated with various neurological dysfunction (Rémi & Noachtar, 2010). These alterations can last many days to weeks, such as psychosis (Krauss & Theodore, 2010; Jackson, 1875) (which in turn has been linked to otherwise unexpected violence (Kanemoto et al., 1999)), or delirium (Krauss & Theodore, 2010); the spatiotemporal patterns studied in the immediate post-ictal period here are unlikely to provide insights to these longer lasting perturbations. Other alterations, however, occur sub-acutely, such as aphasia (Privitera et al., 1991; 1996), paresis (Rolak et al., 1992; Pearce, 1994), cognitive deficits (Helmstaedter et al., 1994; Aldenkamp et al., 1996), and depression (Fisher & Schachter, 2000). Furthermore, in the immediate minutes after seizure termination disturbances to vegetative functions, such as coughing (Hoffmann et al., 2009; Fauser et al., 2004) and hypersalivation (Janszky et al., 2007) can be found.

Indeed, in patients with complex partial seizures, such as studied here, post-ictal alterations have been found in 69 to 81% of seizures (So & Blume, 2010; Kaibara & Blume, 1988; Jan et al., 2001). Beyond this need to further understand the underlying alterations that give rise to these sub-acute symptoms, the early post-ictal activity may well also provide insights into seizure termination mechanisms, with any effects on the recruited neurons' likelihood of firing an action potential likely to still be in place in the immediate aftermath. Furthermore, certain post-ictal symptoms (such as hemiparesis, speech disturbances in cases where the seizure reached the language processing areas, or indeed nose wiping (Leutmezer &

Baumgartner, 2002)) have already been shown important for their ability to lateralize seizure onset.

I reported in the previous chapter that the firing rate of units alone could not account for these post-ictal alterations, because they return to the pre-ictal range shortly after seizure termination (Figure 4.13). Indeed, Fisher & Prince (1977) showed the ability to stimulate neurons to action potential shortly after generalized seizures in a feline penicillin model, suggesting that an inability for neurons to fire is not the cause of the protracted post-ictal state. Instead, it is likely that alterations in the specific firing patterns of the neurons as a population after the seizure are the cause of the alterations found post seizure termination.

### 5.3.1. *Entrained network activity in the post-ictal period*

Qualitatively, evidence for such alterations in firing was found first in the raster plots, showing bursting activity across all units in the post-ictal period in the ictal core, with periods of quiescence between each burst lasting up to a matter of seconds (Figure 5.1). The correlations at zero time lag between firing rates across the whole population in the post-ictal period, with reference to a time matched pre-ictal epoch, provided initial quantification of this alteration to population firing activity (Figure 5.3). The post-ictal correlation coefficients were higher than their pre-ictal counterparts in patient 1, in whom the MEA was within the ictal core during the seizure, and also higher than both time points relative to those from patients 3 and 4, in whom the MEA had remained penumbral.

Importantly, the high correlation between pairs of cells was less dependent upon firing rate than in the pre-ictal period (Figures 5.4 A & 5.5 A). This shows that while some units were indeed likely to have coincident firing given their rates of spiking in the pre-ictal period, the activity in the post-ictal period was highly similar across units of differing firing rates. While the post-ictal increase in correlation was significant in patient 3, the correlation did not occur regardless of firing rate across the whole population, nor was it of the same magnitude. In fact, this penumbral post-ictal increase in correlation was still lower than the pre-ictal correlations in the soon-to-be ictal core, though patient 4 showed similar levels of correlated activity in the pre-ictal period, and so this is unlikely of predictive quality as to future recruitment, but rather brain state, or region, specific.

Of interest, however, is the distinctive grouping of highly correlated pre-ictal firing across patient 1's seizures. When normalized to the firing rate in the pre-ictal epoch, note how in seizures 1 and 3 (Figure 5.5 A (i) & (iii)), the trend shows units with the higher firing rate being more correlated (lower right corner), whereas in seizure 2 (Figure 5.5 A (ii)), the units with lower firing rate displayed the higher correlation (upper right corner). Notably, seizures 1 and 3 in patient 1 occurred at 05:40 and 03:25 respectively, whereas seizure 2 occurred at 21:53. One potential cause of this alteration to the population activity, therefore, is differing brain states. Of interest for further study therefore would be whether this correlation between the more active units is in fact a feature of slow-wave sleep, with spiking well entrained to an underlying delta oscillation in the most active neurons.

### 5.3.2. *Unit specific intrinsic firing alterations*

The next question therefore, was whether neurons were displaying different intrinsic firing patterns, relative to themselves. Calculating the autocorrelation for every unit separately in the pre-ictal and post-ictal epochs showed that the range of different intrinsic properties decreased across the population in the ictal core (note the decrease in standard deviation evident in the post-ictal data in Figure 5.6 A). This reduction in standard deviation in the mean post-ictal autocorrelations within each unit shows that there is a large decrease in the level of "randomness" in the firing times for each unit; knowing when one spike has occurred gives us a higher level of confidence in estimations of when other spikes from the same cell will occur than in the pre-ictal epoch equivalent autocorrelations (Figure 5.6 A). This reduction in randomness in the spike times within each unit suggests that, rather than there being simply high entrainment within the network but with arbitrary bursting, there is instead an external driving force to the network, since there is order to be found within the spike times. Furthermore, any given neuron was found to be unlikely to fire again after a matter of less than a second, up to many seconds later; a feature not found in the penumbral patients. While patient 4 did show an alteration in the intrinsic firing properties of the units, the features were in fact the opposite relationship to patient 1, with the post-ictal period showing less entrained bursts with subsequent quiescence. Seizures were occurring in clusters in this patient, and the post-ictal increase in autocorrelation corresponds to a second seizure within the epoch, although the quiescence prior to onset is of interest.

### 5.3.3. *Population alterations in entrainment during the post-ictal period*

While these autocorrelations indicated that individual units had quite specific patterns of bursts of activity followed by seconds of quiescence, alone, they give no insight to the activity at the population level. Units may be independently undergoing this pattern, with no reference to one another, or alternatively, they may be entrained to one another and bursting as a population. Qualitatively, the raster plots suggested entrainment across the population (Figure 5.1 A), and indeed, supplementing the autocorrelations with the mean cross correlations between every unit pair revealed that the same entrained burst, followed by quiescence, could be found across the whole population, firing as a group at specific moments in time. These bursts were quite evident in raster plots of the spike times across all ictal core seizures (Figure 5.9 B (ii)).

If all units are firing at the same time, then the possibility for carrying information, and thereby performing any meaningful cortical processing, must be severely limited. This was evident in the calculations of mutual information, showing a large post-ictal increase that was not evident in patients 3 and 4. This altered activity of bursts and quiescence was found to last a matter of beyond the first 30 minutes post seizure termination. Note that while the intensity and widths of these bursts were found to return to pre-ictal ranges after just under 30 minutes, this was a calculation of how long until 50% of the events met the pre-ictal range. Furthermore, the range did not calculate whether bursts were still evident, but rather if any bursts present fit the description of the non-bursting period under the same tests. This was done because to quantitatively define the transition from bursting to ongoing activity is not feasible due to it being a slow transition, with bursts slowly becoming less temporally distant, leaving little evidence to go on objectively. As a result, the time delay until physiological processing calculated here is likely an underestimation by some margin.

### 5.3.4. *Single unit spatiotemporal patterns and the protracted post-ictal state*

Perhaps then, this impaired cortical processing through entrained bursting in the post-ictal period, which is not evident in penumbral territories, is related to the sub-acute symptoms of the protracted post-ictal state? Many of the aforementioned symptoms could plausibly arise due to such cortical activity in specific brain regions, disrupting that location's, and its downstream locations', physiological activity. Although this link between the majority of these symptoms and the altered population spike timings found here is unsure, it is still

important to address the effects of these bursts, because, whether or not they directly relate to the multitude of post-ictal symptoms, they are likely to have an effect on cognition. Bursts of activity during the inter-ictal period have already been shown to disrupt cognition in hippocampal recordings from rats (Kleen et al., 2010), and resective surgery has been shown to have the potential to ameliorate this disrupted cognition (Elger et al., 2004). Although these inter-ictal bursts are likely different entities to the post-ictal activity described here, it might be reasonable to anticipate a potentially similar level of disruption to cognition due to the altered network activity found during the post-ictal bursts.

Interestingly, a reduction of cellular excitability has been put forward as a potential cause of both seizure termination and the protracted post-ictal state (Löscher & Köhling, 2010). Intracellular $Ca^{2+}$ concentrations are known to rise throughout seizures, thereby activating $K^{+}$ channels (Timofeev et al., 2004), which would decrease the input resistance. Indeed, input resistance in granule cells has been shown to decrease in acute slices from temporal lobe epilepsy surgeries (Stegen et al., 2009), which would result in greater synaptic input being required for the cell to reach threshold. The post-ictal bursts of action potentials found in the cortex that was recently fully recruited to the seizure is in keeping with this hypothesis; if all cells have increased leak conductance then they would all remain sub-threshold and quiescent until a very high level of input is reached, which would result in a majority of cells firing within a small window, in bursts.

Alternatively, alterations to synaptic connectivity may give rise to these ongoing, entrained bursts. Indeed, reduction of cellular excitability due to ionic concentrations could be expected to return to baseline levels sooner than the return to baseline-like activity found here, unlike the return to baseline of any resultant changes in synaptic connectivity. In fact, cell assemblies that fired during seizures have been shown more likely to be re-activated during post-ictal slow wave sleep (Bower et al., 2015), thereby providing tentative evidence for a hijacking of the consolidation mechanism proposed for memory formation (McGaugh, 2000).

Many other reasons have been proposed for the protracted post-ictal state however, to which this altered firing activity would have varying degrees of applicability. Opioid receptor upregulation has been posited as a likely cause, both of seizure termination and the post-ictal state, and at a similar timescale of at least a matter of hours to return to baseline (Hammers et al., 2007). Extracellular acidification increases during a seizure (Lux et al., 1986),

and has been shown capable of stopping seizures (Xiong et al., 2000; Caspers & Speckmann, 1972). In fact, alkalosis has been shown to precede febrile seizures (Schuchmann et al., 2006), however the pH of the extracellular environment is likely to return to baseline at a much quicker rate than the return of the firing patterns shown here. The ionic environment has also been suggested in seizure termination in acute *in vitro* models (Bikson et al., 2003; Pinto et al., 2005), though the firing patterns found here also outlast the expected duration of alterations in ionic concentrations. Cerebral blood flow has been shown to increase in the post-ictal period (Weinand et al., 1994), normalizing to baseline levels within half an hour (Berkovic, 2000), which would plausibly be consistent with the time scale of the bursting patterns presented here. That said, it is important to bear in mind that the time scale for the return to baseline activity of the post-ictal bursts found here is a measure of how long until detected bursts are statistically indistinguishable from a period during which there are no bursts, at a probability level of less than 0.01. As such, this delay suggests a time point at which it is difficult to differentiate between true bursts and chance detections amongst physiological neuronal activity, rather than a delay until the return of truly physiological activity. It is likely, therefore, that the delay until baseline spatiotemporal activity patterns outlasts the probabilities of detecting differences with the methods reported here.

Overall, it is likely that the array of post-ictal symptoms mentioned arise from a mixture of the above mechanisms, and whether or not this characterised firing activity found in the post-ictal period relates to the symptoms is, as-of-yet, undetermined. Due to the subdural electrode analyses being performed upon burst-triggered averages, and there being no bursts in patients 3 and 4, equivalent control analyses could not be performed for the penumbral territory, using these methods.

Regardless of their aetiology, which is discussed further in the following section, the downstream effect of these post-ictal bursts is a potential cause for comorbidities often found in epilepsy. Furthermore, long-term deficits in cognition have been found in rats as a result of inter-ictal events during development (Khan et al., 2010), and so, whilst not directly equivalent to inter-ictal activity, the bursts found during the post-ictal period here may be of special consideration in juvenile forms of epilepsy, in whom preventing these bursts of activity may ameliorate ongoing cognitive deficits later in life.

### 5.3.5. *The potential of post-ictal bursting activity as a seizure onset localization tool*

The post-ictal bursts may have strong potential as a method to aid in the localization of the seizure onset zone in resective surgery cases. The highly characterized unit activity is likely also to be evident in lower frequencies, given its distinctive alteration between large amounts of neuronal activity, followed by quiescence. If so, it could play a critical role in aiding the localization of the seizure onset zone. However, it was found that the burst-triggered averages were similar across many of the standard subdural electrodes, implying similar activity regardless of local recruitment to the ictal core. It therefore would be necessary to further understand their mechanism across different regions to accurately utilize possible information within these bursts. This is in keeping with other potential localization tools, such as HFOs, which have been proposed as a localization tool in focal neocortical epilepsy (Blanco et al., 2011), but in fact are likely evident in both the ictal core and penumbra (Schevon et al., 2009). In the case of the HFOs occurring in both regions, the level of phase locking present in the high gamma has been proposed as a method by which to determine the two regions (Weiss et al., 2013; 2015).

In keeping with the alterations found by Weiss et al., the underlying average changes in power through the post-ictal bursts did display characteristic differences near to the ictal core, which were well maintained from seizure to seizure in this patient. This indicates that an analysis of concomitant increases and decreases in power between the high and low frequencies in the standard subdural electrodes may well be a biomarker for the location of such bursting in the post-ictal period, due to the intense spiking riding on the increase in power of the low frequencies.

The high gamma frequencies are typically considered as a correlate for multi unit activity (Ray & Maunsell, 2011), and so at first pass, the burst triggered average oscillations across the subdural array presented here would imply that a burst of multi unit activity is happening concurrently across the whole temporal lobe, if not further afield (Figure 5.15). However, upon inspection of the burst triggered average power across the different frequency bands, an important point arises. Counter intuitively to the increase in high gamma oscillation seen in the burst triggered averages, the power in the high gamma band can be seen to decrease during the burst (with similar results in the lower frequencies also). This is indicative that rather than there being an increase in the activity at these frequencies

at the time of the burst, there is instead a decrease in activity, but an increase in entrainment between events relative to the bursts.

This possible reason for the increase in the average oscillation with a decrease in the average power is explained in Figure 5.18. Note how, although the sine waves in Figure 5.18 A have greater power during the early and late sections, the average across the two examples (Figure 5.18 B) results in no amplitude during those periods, as they are out of phase to one another. Conversely, during the central section, the sine waves actually have lower amplitude in both events, but because they are in phase, the average across them appears to show higher amplitude at this frequency than during the early and late sections. The mean of the powers of the two signals, however, shows the anticipated drop due to the decreased amplitudes during the central section (Figure 5.18 C), despite the concomitant increase in power of the mean signal (Figure 5.18 B). Note that this also highlights the importance of calculating the mean of the powers rather than the power of the mean signal, when analysing average signals.

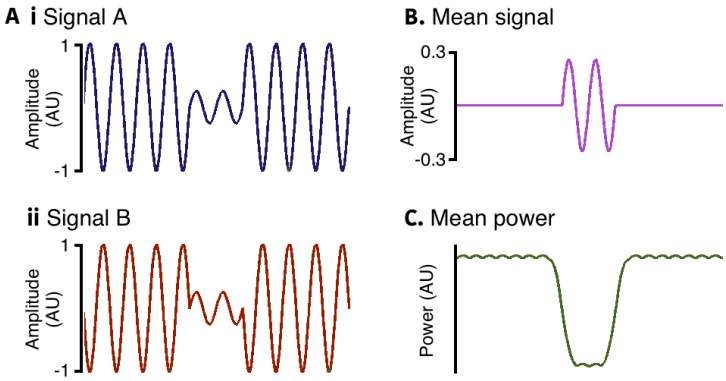


**Figure 5.18. Explanation for the increase in the burst triggered signal despite concomitant decrease in the burst triggered power.**
**A** (**i** & **ii**). Two example signals, both sine waves, though with large amplitude, out of phase, periods during the early and late phase, and with a lower amplitude, but in phase period in the central section. **B.** The mean of the two signals in A, showing a seeming increase in activity at the frequency of the sine waves during the central section, despite the decrease in amplitude of both signals in A, due to the out of phase periods cancelling out. **C.** Decrease in mean power across the two signals, at the frequency of the sine waves in A, during the low amplitude central section. This decrease is expected, despite the seeming increase in power found in the mean signal in B. Note that this highlights the importance of calculating the mean of the power, rather than the power of the mean.

This possible explanation for the discrepancy between the burst triggered average oscillations and powers would suggest that each event is resulting in similar downstream effects across the cortex, but that only in the recent ictal core is there intense, entrained

bursting, leaving the increase in high gamma power solely near the recruited territory (electrodes G60 and G52). This would indicate that in the more distal electrodes, the burst-triggered averages are picking up distributed synaptic activity from the event. Perhaps then, not only would this discrepancy potentially help differentiate between territories that had and had not been recruited during the seizure, but also, an analysis of phase differences and relative delays across the subdural electrodes might provide insights to the seizure trajectory between areas.

Since the distributed synaptic activity was found across a wide area of cortex, this raises the question of where the presynaptic activity originates. It is likely that the synaptic activity across the cortex is a result of subcortical activity. Devergnas et al. (2012) found subcortical entrainment to the cortical ictal activity, in the subthalamic nucleus, globus pallidus and putamen in a macaque penicillin model of focal motor seizures. A likely explanation for the post-ictal bursts, therefore, is that thalamo-cortical loops have resulted in ictal activity progressing to subcortical regions. In turn, these subcortical regions are providing low frequency input to the cortex in the post-ictal period, which results in activity occurring in the ictal core only at the height of said input due to altered excitability.

Indeed, post-ictal thalamic activity has been shown to cause the cortical slow waves found after termination of electrically induced hippocampal seizures in rats (Gummadavelli et al., 2014). Gummadavelli et al. also found that thalamic stimulation during the post-ictal period resulted in the return of exploratory behaviours in the post-ictal period, along with a decrease in cortical slow waves and an increase in multi unit activity. Alternatively, hippocampal theta is known to entrain cortical gamma (Sirota et al., 2008), and so altered activity in the mesial temporal lobe may result in such distributed bursting activity as seen here. Further spatiotemporal analyses of these bursts, incorporating the time lags between their occurrences and their relative component frequencies, are necessary to elucidate the mechanisms at play during this activity.

This intense bursting population activity may, alternatively, purely be a result of travelling waves of activity in the lower frequencies, resulting in these coincident bursts of unit firing, though the marked similarity across cortical regions from seizure to seizure lessens this likelihood. Importantly, it is also worth bearing in mind that brain state plays an important role in neuronal spiking activity. Steriade, McCormick, et al. (1993) showed that sleep is characterized by synchronized bursts of activity in many neurons, resulting in the low

frequency oscillations found in slow wave sleep, and so this is a potential source of the bursts in the post-ictal period. The bursts found here, however, are not characteristic of any physiological sleep stage, and the inter burst interval is initially far longer than would result in the low frequency delta oscillations mentioned. Furthermore, the steady increase in frequency of these bursts is uncharacteristic of sleep related oscillations. Studies of spiking activity during sleep in humans have not reported on spatiotemporal dynamics as seen here, from either depth electrodes (Le Van Quyen et al., 2010), or similar MEAs to those used here (Peyrache et al., 2012).

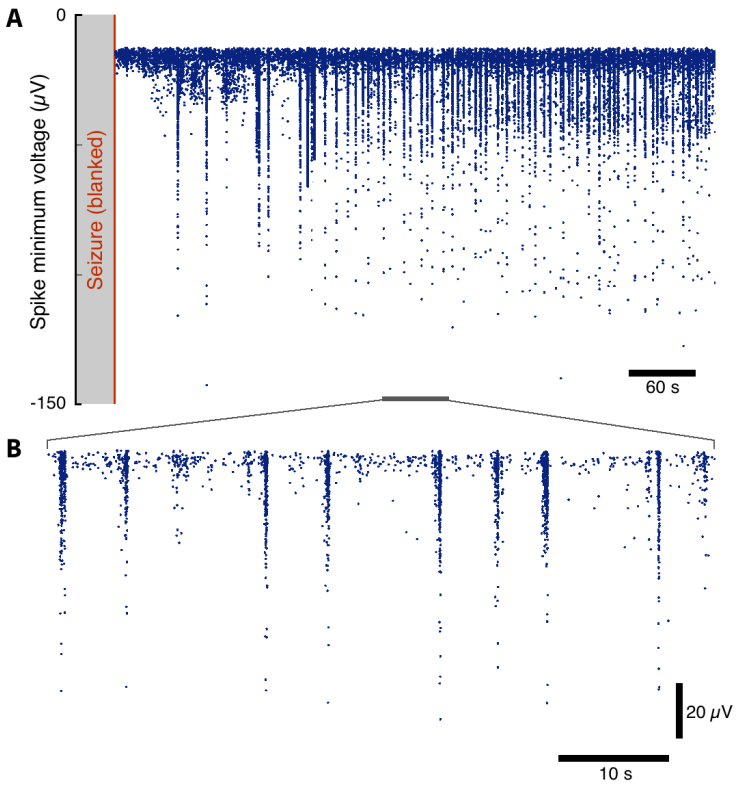


**Figure 5.19. Evidence of post-ictal bursts in the multi unit activity.**
**A. T**he voltage at the detected trough from all spikes, across all electrodes, plotted against time, over 10 minutes from the post ictal epoch in patient 1, seizure 1. **B.** Magnification of a one minute epoch marked by the grey bar in A. Note the readily apparent bursts of activity in the multi unit detected spikes, without the need for separating the spikes into single units.

However, brain state has been shown to alter firing patterns beyond the oscillation specific fluctuations. The cortical state has been shown to determine the single unit response in primary visual cortex relative to the upstream activity of the lateral geniculate nucleus (Schölvinck et al., 2015), and so analyses of differing activity patterns in single unit responses

need to be performed with reference to brain state. To rule out sleep specific brain state as the cause of these bursts, long term analysis of spike timings needs to be performed, in tandem with polysomnographic sleep scoring. This is likely easier than initially anticipated, as the multi unit activity alone, without the need for spike sorting, can identify the bursts found in the post-ictal activity here, assuming the voltage at maximal negative deflection detected is stored also (Figure 5.19). This would be a worthwhile next step to further assess the specificity of this spatiotemporal pattern to post-ictal non-physiological activity.

Regardless, this is ongoing work, and data for the ictal core have thus far been presented only from a single patient, with patient 2's recording having been interrupted during the seizure. As a result, this discrepancy between the ictal core and penumbra may be a patient specific one, rather than a distinction between ictal core recruited and penumbral territories, especially as the bursts are evident beyond the local territory to the MEA. Preliminary data from a second patient in whom post-ictal data from an MEA placed in the ictal core were available, however, shows evidence of very similar single unit and population features, albeit from a shorter duration thus far (Figure 5.20).

Given, however, the evidence of post-ictal bursting on subdural electrodes distal to the MEA in patient 1, it seems likely that this activity is not specific to the region of recruitment to the ictal core, but rather to a subset of seizures. Preliminary observations by the treating epileptologists suggest that these bursts are related to semiology seen in the concurrent video recordings, and that these events occur predominantly after generalized seizures, and to a lesser extent, after complex partial seizures as analysed here. This does not, however, discount the potential seizure onset localization capability that may be provided by this post-ictal activity in seizures where it occurs, and further analyses are called for. Furthermore, these analyses have focused almost solely on the spiking activity across the isolatable units during the post-ictal period, with only preliminary comparisons to the lower frequency data performed thus far. These data will benefit considerably from further insights into the oscillatory activity occurring during the same post-ictal period. At present, these equivalent studies are ongoing at Columbia University on the same dataset, and will likely further elucidate the mechanisms underlying this bursting activity when combined with these single unit findings.

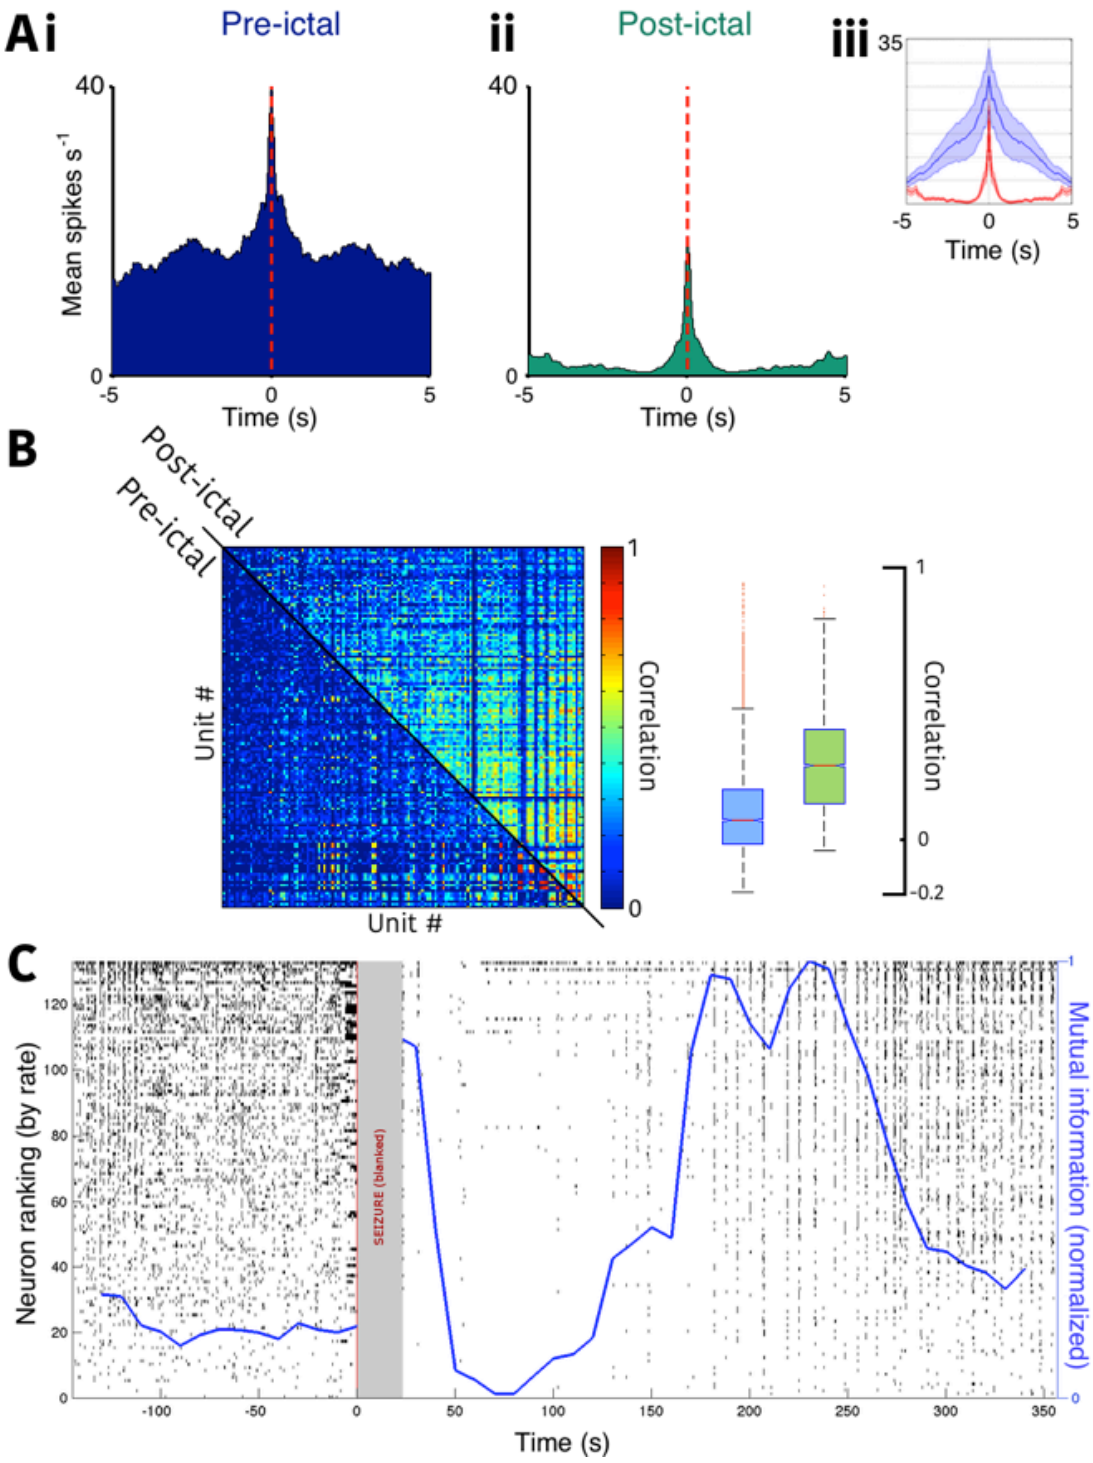


**Figure 5.20. Equivalent population bursting activity in second patient from ictal core territory.**
Preliminary analyses of recordings from a 5th patient, in whom the MEA was also within the ictal core territory during seizures, showing the same post ictal bursting as in patient 1's seizures, suggestive that this activity is not solely a patient specific post ictal response. **A.** (**i**)**.** Pre-ictal and (**ii**) post ictal mean unit cross correlations, as in Figure 5.7, with (**iii**) the mean autocorrelations from the pre and post ictal period, as in Figure 5.6, all showing the same hallmark patterns of the three seizures in patient 1. **B.** Post ictal increase in correlation in firing between all possible unit pairs, with the post ictal epoch in the upper right triangle of the matrix, and the pre-ictal epoch in the lower left. These results correspond to 8,778 possible unique pairs, from 133 single units. **C.** Raster plot as per Figure 5.9 A, ordered by single unit firing rate over the full epoch, from 3 minutes prior to seizure onset, until 6 minutes post. The seizure has been blanked. Mutual information, normalized to its own maximal value, is in blue. Note the post ictal increase in mutual information, after a brief period of low mutual information due to quiescence in the majority of units.

# Chapter 6. General Discussion

The previous three chapters each closed with specific discussion regarding their findings, and so I close now with a more general discussion regarding the issues and utility of these methodologies, along with brief comment on potential future directions.

## 6.1. The issues of rare recordings in humans

Foremost, the limitations due to the sample size of the dataset analysed here must be acknowledged. Invasive recordings in humans are necessarily rare. Furthermore, it is hard to control for variability in these data; the recordings that are possible in humans are always dictated by clinical reasons, and thus equivalent control recordings in humans are rarely available. In animal studies, breeding allows for examination of homogeneous populations, and environments can be maintained from animal to animal. Neither of these situations is plausible in humans. As a result of their infrequency, there is a necessity to pool recordings from humans that have similarities, but are not necessarily equivalent. These various considerations all serve to increase potential noise in the data.

As a result, it would seem that little insight might be gleaned from these MEA recordings in humans, at least at the current number of patients; each individual case is a single biological replicate. It may not be plausible to extrapolate results from a single patient to seizures in general, though if so, this highlights an important point – seizure aetiology may be diverse, and a future direction of therapy is likely through patient specific, tailored treatments. That said, it has been argued that the mechanisms of seizure propagation and termination are likely similar across patients regardless of the cause of ictogenesis, through a common dynamical mechanism (Kramer et al., 2012), and thus extrapolation from the recordings across these patients would be relevant beyond solely these cases. Indeed, it seems likely that whilst the focal ictogenic origin may show a great variety of pathology, when the ictal activity propagates into more physiological territory there is likely to be a convergence of

activity patterns. Therefore, the seizure termination mechanism may be less variable than those of seizure initiation.

I would also draw parallels between the issue of rare recordings, and that of rare diseases. Study of single patients is common in human genetics, and it has been suggested that reliance on these results is unwise due to the lack of statistical power (MacArthur et al., 2014). However, Casanova et al. (2014) put forward the case that these results from single patients can indeed be conclusive, provided the data are analysed with reference to rigorous experimental validation *in vitro*. In fact, Casanova et al. report that 21% of single gene human primary immunodeficiencies were originally reported in single patients.

With reference to the rarity of the recordings studied here, rather than the rare diseases studied in the aforementioned studies, there is an equivalent situation. There is a wealth of data originating from animal models of epileptiform activity, and the power of these single patient recordings lies in part in their ability to test hypotheses arising from these animal studies, and indeed, vice versa. The assumption is that human seizures are equivalent to these ictal recordings in animals, be it *in vivo* or *in vitro*. There are definite similarities in cellular and cortical structure between humans and the animal models, but in the finer details it is unknown how relevant these models might be.

To derive hypotheses on the animal recordings and then test them on data such as reported here allows for powerful insights even in the small dataset; it takes very few recordings to disprove a model. Of course, this is not without its caveats. In the recordings studied here, it is important to take reference to which cortical region the data are being collected from – an absence of proof is not necessarily a proof of absence. This realization of the need to take reference of the cortical region, however, illustrates the power of the combination of animal work and single patient analyses. The animal studies provided an important context to help make sense of these complex recordings, and it is from the animal work that the hypothesis of the ictal core and penumbral territories were derived. An analogous relationship is found in the ability to interpret these rare recordings relative to more common techniques in humans. For instance, Weiss et al. (2015), performed analyses on the more widely used ECoG recordings, as an assessment of findings from the MEA data (Weiss et al., 2013). Combined, these stress the importance of the utility of these rare recordings; they allow for direct testing of hypotheses formed in animal models, or in the more widely available recordings from humans, and similarly, they allow for the formation of hypotheses on data

that would otherwise be unavailable that can then be tested using the more widely available techniques.

The assessment of the validity of animal models to human seizures, as is available in these data, is vital. There is a vast array of animal models to accompany the wide range of seizure types, and it is not rare for these models to be chosen due to convenience or tradition, rather than a thorough assessment of their applicability to the hypothesis being tested (Grone & Baraban, 2015). In fact, it is perhaps partly a result of a continued reliance on potentially ill-suited animal models that new anti-epileptic drugs have shown little improvement to the percentage of pharmaco-resistant cases over the last two decades (Löscher, 2011; Bialer & White, 2010).

That is not to suggest that animal models should be cast aside, but rather that an ideal test of the hypotheses derived from these models lies in the rare recordings from humans. There are many experimental techniques that are simply unfeasible in humans. Action potentials involve a concomitant increase in intracellular calcium ion concentration (Baker et al., 1971; Jaffe et al., 1992; Svoboda et al., 1997), and so calcium imaging can provide considerable information of network activity with both good temporal and spatial precision (Peron et al., 2015). However, this requires either the introduction into the cells of fluorescent dyes that chelate calcium ions, or the transfection of genetically encoded calcium indicators (GECIs), neither of which are feasible at present in humans.

Similarly, as discussed primarily in chapter 4, it is only with recordings from confirmed single cells that we will have definitive proof of the PDS in human seizures, and the necessary technique for such recordings, patch clamping, is also not feasible in humans. A promising future direction for the control of seizures lies in a closed loop system utilizing optogenetics, allowing for neuronal intervention upon the detection of ictal activity. This has been shown effective in animal models (Tønnesen et al., 2009; Wykes et al., 2012; Paz et al., 2012; Krook-Magnuson et al., 2013; Wykes et al., 2015; Kros et al., 2015), and a roadmap to its use in humans has been proposed (Kullmann et al., 2014). The future of optogenetics in humans will therefore provide another potential source for the validation of these animal models, but will remain, sensibly, dictated by the clinical needs, leaving the animal models the best source on which to form hypotheses, to be tested *in vivo* in humans.

*In vitro* studies of human brain slices from resective surgery are an underutilized middle ground (Jones et al., 2015). Whilst still considerably less numerous than data from animal models, resections are, of course, more common than MEA implants in humans, and they allow the use of the aforementioned techniques that are not feasible *in vivo*. Not only do these bypass the use of animals, but they can also allow for analyses of spontaneously occurring inter-ictal events (Huberfeld et al., 2007; Roopun et al., 2010; Cunningham et al., 2012; Simon et al., 2013), thereby bypassing the potential confounds found through the choice of pharmacological model in the equivalent animal models. That said, spontaneously occurring full ictal-like events in these recordings are still rare (Jones et al., 2015; Cunningham et al., 2012), and thus this remains an issue that these MEA recordings will be able to supplement.

I would also suggest that an understudied mechanism within seizures is that of termination, and that the *in vitro* brain slice preparation is perhaps partially unsuited to this particular mechanism. As mentioned, the choice of pharmacological model will have bearing on many aspects of the epileptiform activity, including that of termination, due to whether they elicit pathological activity through increased excitability of the network, or decreased inhibition. Further to this issue, I would posit that the relative size of the network available in the *in vitro* slice results in a more rapid loss of unrecruited territory, especially if combined with a disinhibition model, which would result in a much faster speed of propagation (Traub & Miles, 1991; Chagnac-Amitai & Connors, 1989; Pinto et al., 2005; Albowitz & Kuhnt, 1995; Wadman & Gutnick, 1993). This reduction in "viable", unrecruited tissue for the ictal activity to propagate into would likely alter the dynamics of termination of this activity when contrasted with intact cortical and subcortical structures *in vivo*.

For these reasons, I suggest that whilst the statistical power from this low number of replicates would initially appear low, these recordings can be vital to our understanding of seizure mechanisms, assuming the results are considered with careful reference to the statistically powerful observations found in more widely available recordings. Whether or not analyses of the spatiotemporal patterns of single units (which represent a relatively low sampling when compared to the total number of neurons in the region recorded from) are a good source for the elucidation of these mechanisms is, however, not certain. I therefore address this question in the following section.

## 6.2. The utility of extracellular action potential recordings

In chapter 3, I argued that it is with recordings of action potentials from single units that the most insight will be gained into the links between consciousness and cortical processing. Whilst a holistic approach, incorporating single unit recordings, would of course provide the best platform for answering these questions, ultimately, in the majority of cortical processing, neural computation is performed at the level of action potentials and their specific timings. It is due to this that I suggest unit recordings may provide significant insights when using the right analytical tools.

Although we are still unsure of the exact origin of the local field potential (LFP; Łęski et al., 2013), it is generally accepted that it corresponds to a superposition of all ionic process ongoing within the area (Buzsàki et al., 2012). The majority of these fluctuations however only result in alterations at the processing level via an adjustment to the probability of a neuron reaching the threshold for action potential at any given time. To clarify, that is not to say that oscillations are not important for functional processing, but that it is in the resultant adjustments to spike times that the oscillations provide that computation occurs, and thus ultimately, the spike times themselves are the level at which cortical processing occurs.

Whether or not gamma oscillations are truly representative of cortical processing has been called into question, however (Merker, 2013). Notably, the power of gamma band frequencies increasing during post-ictal unconsciousness relative to baseline, conscious activity in human subdural electrode array recordings (Pockett & Holmes, 2009). It has therefore been posited that gamma oscillations are, in fact, merely the “sound” of active cortex rather than being a functional component of cognition, arising from each neuron having resonant frequencies that combine to form oscillations during activity (Merker, 2013). While the power in the gamma band is not directly related to the activity of any individual unit, it is tightly related to the level of correlation between a population of neurons (Nir et al., 2007). It is likely that gamma oscillations arise through a multitude of mechanisms (Whittington et al., 2010), however, and oscillations at different frequencies within the gamma range could have differing roles in cortical processing (Sedley & Cunningham, 2013). Differing roles of oscillations within gamma are in fact evident, with dual frequencies, generated in separate locations, allowing for specific control between the layers of cortex (Ainsworth et al., 2011).

Furthermore, gamma oscillations have since been shown capable of adjusting perception. Helfrich et al. (2014) used high density transcranial alternating current stimulation to induce synchrony, in the gamma band, between the two hemispheres in primary visual cortex in humans, and in doing so, were able to manipulate visual perception in humans. Moreover, Lakatos et al. (2008) demonstrated that the low frequency delta oscillations in macaque primary visual cortex become entrained to the rhythm of an input stimulus, thereby playing an important part in functional processing. This same entrainment to input streams of stimuli has since been shown also in human subdural neocortical recordings, showing that attentional selection is mediated by entrainment of these low frequency oscillations, but in order to adjust spiking probabilities to become in line with the input (Besle et al., 2011).

Theta oscillations, too, have been shown relevant to functional processing by providing a "computational window" during which spiking may occur, such that multiple local circuits in the entorhinal cortex and hippocampal loop may perform independent processing that can still be translated from one area to another (Mizuseki et al., 2009). Moreover, the timings of spikes in hippocampal place cells, relative to an underlying theta oscillation, have been shown to encode position, with this phase precession providing higher accuracy to the computation than that of the firing rate of the cells alone (O'Keefe & Recce, 1993). Furthermore, the strength of memory recollection in humans can be predicted by the level of phase locking between spike timings and an underlying theta oscillation (Rutishauser et al., 2010). All of the above oscillatory adjustments to cortical processing, however, is through the modification of the probability of cells reaching action potential threshold within a certain window of time, with the single neuron spike timing being controlled by the oscillations (Jacobs et al., 2007). As a result, it is at the level of single unit spike times relative to one another that cortical processing is likely to occur.

Furthermore, the oscillations themselves are predominantly a downstream result of single unit activity, with many frequencies thought to be mediated by inhibitory interneurons (Buzsàki & Watson, 2012; Buzsàki, 2006; Whittington et al., 2000). Delta oscillations too, are the result of spike timings in populations of neurons, with ON and OFF periods. Increases in firing rate occur during the ON period, with a following quiescence of suppressed firing in the OFF period, resulting in slow wave oscillations (Csercsa et al., 2010). However, gamma oscillations in the CA1 region of the hippocampus have been shown to occur without the necessity for the input of the fast spiking interneurons (Craig & McBain, 2015), and there is

evidence for non-fast spiking interneuron induced gamma activity in the neocortex also (Adesnik & Scanziani, 2010). As such, the relationship between the lower frequency LFP and spiking activity of neurons is still not known definitively.

Ultimately however, the functional processing of the oscillatory activity in the lower frequencies occurs through the resultant membrane potential fluctuations in individual cells. These fluctuations can adjust the probability of that cell reaching the threshold for an action potential at any given time (Engel et al., 2001), and it is whether or not the action potentials occur at that time that provides functional processing. Therefore, with so little yet understood about seizure onset, propagation and termination mechanisms, and given that seizures are a disorder of aberrant neuronal firing, the insights from single units during these time points will be invaluable, with two major caveats.

The first caveat is that, while it is eventually the activity of these units that provides the outward manifestations of a seizure, the aetiology may not be contained to those cells; for example, a glial contribution towards temporal lobe epilepsy has been demonstrated through astrocytic control of neuronal NMDA receptors (Tian et al., 2005; Clasadonte et al., 2013). As a result, the optimal therapeutic target may well be upstream of any aberrant alterations in activity found in the single units. The second caveat, and most importantly regarding the interpretation of these types of recordings, is that the downstream effects of a given unit must be known in order to truly understand the resultant network activity of a spike occurring. As it stands, I reported in section 3.2.4 that subclassification of single units into putative cell types did not provide a clean bimodal distribution on any metrics, leaving it unwise to draw conclusions regarding cell type by such methods in these data. These methods have been shown to result in the expected firing rates of the fast spiking and regular spiking populations in neocortical recordings in both macaques (Mitchell et al., 2007) and humans (Peyrache et al., 2012). Therefore, assuming that the method is plausible (despite the evidence that the main classifier of fast spiking interneurons – a short spike half width – is known to not be specific to that cell type (Vigneswaran et al., 2011)), would these separations provide the insights necessary to assess network function?

I would suggest that even where these separations prove to be a reliable indicator between fast and regular spiking cells, such a subclassification does not result in enough information to tease apart the network function fully. The resultant downstream network activity varies

considerably by different interneuron types alone. Firstly, the resultant changes in activity in the postsynaptic cell is dependent largely upon which postsynaptic neuronal population it targets (Pfeffer et al., 2013), and on those cells, whether it targets the peripheral or proximal dendrites, soma or axon initial segment (Wilson et al., 2012; Lee et al., 2012; Atallah et al., 2012; Pouille et al., 2013; Freund & Buzsàki, 1996; Somogyi & Klausberger, 2005; Isaacson & Scanziani, 2011). Secondly, while a general "blanket of inhibition" upon the excitatory pyramidal cells is provided by most inhibitory cell types over distinct time periods, the vasoactive intestinal protein containing interneurons (VIP) instead remove that inhibition locally, by inhibiting the somatostatin, and to an extent the parvalbumin, containing interneurons, thereby disinhibiting the rest of the network (Lee et al., 2013; Pi et al., 2013; Karnani et al., 2014). As such, merely separating single units into the standard fast spiking and regular spiking groups does not provide knowledge of the postsynaptic effects from that unit.

Given these caveats, studies of population single unit activity with reference to the more readily accessible data are invaluable towards our understanding of the relationship between oscillatory activity and the underlying neuronal processing. In the future, however, once the associations between the more readily accessible oscillation data and the underlying neuronal activity are understood with greater specificity, the reduced benefit of these invasive procedures is likely to become outweighed by the cost. The future of such recordings is discussed in the following section.

## 6.3. Future studies using similar devices in pre-resective surgery epilepsy patients

The era of recording single unit spikes in humans may be a brief one, on-going only whilst we need to make better sense of the data available from less invasive recordings. I discussed in section 6.1 how it is with these data that the hypotheses derived from animal models may be ultimately tested, though once we have a better understanding of how activity in EEG, ECoG, functional magnetic resonance imaging (fMRI) and magnetoencephelography (MEG) all relate to each other, and to the relevant neuronal activity during these events, then these hypotheses will be testable using those, less invasive, and more widely used technologies. Therefore, perhaps the most important aspect of ongoing studies such as these, is to relate these activities, classifying spatiotemporal activity patterns that are specific to the various

ictal and inter-ictal time points and regions, and assessing average lower frequency activity across many examples. It will be important to sub-classify different oscillations and to link them to brain state and specific cortical processing mechanisms found in the spatiotemporal activity of the single units.

It is likely that in doing so, we may also gain insight into previously discarded, readily accessible information, such as the DC shift (Ikeda et al., 1996). Currently, many EEG amplifiers do not provide the facility for recording the DC shift, yet it has been suggested to be ubiquitous (Ikeda et al., 1999; Kim et al., 2009; Rampp & Stefan, 2012; Wu et al., 2014). To assess the exact time of the DC shift at different electrodes, with reference to the time of recruitment to the ictal core would be worthwhile, though it is notable that in the recordings described in this thesis, only one patient showed evidence of a DC shift. This singular occurrence most likely arose due to these slow changes being filtered out by the 0.3 Hz high pass filter in the hardware, and so adjustments to the recording set up would be necessary to perform this analysis. Beyond this, the DC shift has been shown to be relatively cortically widespread (Kim et al., 2009), and so potentially may highlight the entire ictogenic network (Wendling et al., 2010). It remains plausible that the wavefront may be localized through the timings of the DC shift (Thompson et al., 2015), and more recordings are necessary to address this possibility.

In the immediate future, given the distinction between the ictal core and penumbral territories during seizures through multiple features (Schevon et al., 2012; Trevelyan & Schevon, 2013; Weiss et al., 2013; Merricks et al., 2015), it would be pertinent to make the same distinction when analysing recordings of epileptiform activity in humans. This is important not only during the peri-ictal time point, but also during inter-ictal periods. For example, heterogeneous firing has been found to occur during epileptiform inter-ictal discharges (IIDs) from equivalent MEA recordings in humans (Keller et al., 2010), but without explicit reference to whether or not that cortical region was subsequently fully recruited to the ictal core. In fact, it is likely that IIDs follow similar recruitment patterns to those of the seizure, with downstream effects occurring within the penumbra, as has been suggested in high frequency oscillations with a secondary response beyond the core (Schevon et al., 2009). If this were the case, it could be expected that the downstream recruitment from an IID that occurred within the seizure onset zone would be heterogeneous, in keeping with the

heterogeneous response found in the penumbral territory during the seizure, as reported here, and by others (Truccolo et al., 2011).

Notably, however, analysis of IIDs of an equivalent heterogeneous nature (and therefore, I suggest, also from penumbral territory) from hippocampal or entorhinal cortex depth electrode recordings in humans, found a significant alteration in unit activity some hundreds of milliseconds prior to the IID (Alvarado-Rojas et al., 2013). This alteration was also heterogeneous, and in the aforementioned study they conclude that this suggests IIDs arise not from hypersynchronous bursts, but instead are generated by a sparse neuronal network of pre-IID synchrony. An alternative interpretation might be that these events are a continuous downstream response to an IID in the ictal core. More analysis needs to be done to help decipher these possibilities, with explicit reference to the ictal core and penumbra; for example, one might expect that during IIDs in the subsequent ictal core, the PDS would give rise to waveform changes similar to during the seizure. Even if the change were minute, such alterations in spike shape are exaggerated in the extracellular space and would likely result in loss of distinct clusters, as has been shown during physiological bursts of activity (Henze et al., 2000; Harris et al., 2000). As such, prior to discussing significant increases or decreases in spiking during any epileptiform event, it is necessary to be certain that these are genuine alterations to the number of action potentials from the cell, and not simply alterations to the waveforms that cause the spikes to not be attributed to that unit.

Importantly, this necessity to take reference to the exact pathological location of the recordings is also the case when considering the potential of seizure prediction. The ability to provide warning of an imminent seizure, without a high level of false positives, would substantially improve the patient's quality of life. Furthermore, closed loop control of seizures will ideally make use of reliable early detection, if not prediction, to avoid side effects during inter-ictal time points from false positives, or worse, failure to intervene during a seizure in the event of a false negative.

It seems likely, given the large level of low frequency activity in the penumbra whilst the ictal focus is elsewhere, that recordings from different parts of the brain will have different predictive value. Work into potential methods of predicting seizures has been ongoing for nearly four decades, with results tentatively indicating worthwhile progress (Gadhoumi et al., 2015). In their review, Gadhoumi et al. provide an overview of a variety of seizure prediction

methods, with variable success rates. As they point out, future work on effective seizure prediction and detection will benefit considerably from the increasing number of databases of EEG recordings. I suggest that on top of this, post hoc assessment of the pathological location of these recordings, given the new ability to define them in the lower frequencies (Weiss et al., 2013; 2015), will provide considerable benefit to the work on seizure prediction.

I would close with a brief thought on the existence of these invaluable databases of EEG, and especially MEA, recordings. Given the scarcity of these data, and the issues outlined in section 6.1, an important future direction for this research is to progress towards an open model of data sharing. As it stands, with the current system of scientific publication in medical research, the issues of authorship impede the creation of these open datasets. It is important to progress towards a higher level of collaboration with regards to these studies, as is becoming the case in other areas of science, such as the work in high-energy physics from the Large Hadron Collider at CERN (Aad et al., 2015). This, of course, provides issues of accurate credit (Strange, 2008), and thus a considerable alteration to the scientific publication model may be required as a result. If the cost of the current system, however, is a lack of reproducibility in results or an unnecessary limitation to statistical power, I would suggest that the benefits of an open data system likely outweigh the costs of such a change.

## 6.4. Conclusions

Continuous stability of single units recorded in MEAs in humans during the subacute time frame has not been shown before. In this thesis, I have demonstrated that recordings performed during this time period in pre-resective surgery epilepsy patients allows for single units to be accurately followed over many hours to days. This therefore allows for insights into the spatiotemporal dynamics underlying various epileptiform activities, including seizure onset, propagation and termination, at an individual neuronal level in humans.

This stability was used as a baseline to show how single unit alterations occur during seizures, in specific cortical areas. This further confirmed the hypothesis that seizure activity can be subdivided into the core, recruited cortex, and the surrounding penumbral region, finding that while the penumbral response to the intense synaptic barrage resulted in

heterogeneous activity, the ictal core instead showed intense, hypersynchronous firing, with evidence of paroxysmal depolarising shifts. The penumbral response provides evidence that propagation of ictal activity is constrained at the wavefront through feedforward inhibition, thereby withstanding the intense glutamatergic barrage from the ictal core.

Subsequently, the same spike sorting showed that even in the ictal core, firing rate returns to baseline levels in a matter of seconds to minutes, and thus the protracted post-ictal state must arise either at a population level, or through recruitment of subcortical structures. Given this hypothesis, I provided preliminary evidence for previously unseen patterns of neuronal firing in a subset of seizures after termination. These findings lay the groundwork for future single unit analyses of peri-ictal activity, and confirm that long-standing features of ictal activity in animal models are applicable to human epilepsy also.

# Appendix: Related contributions to other studies

During the course of these studies, I also helped devise methods for analysing multi-unit activity as recorded from tetrodes in mouse brain slices *in vitro*, as described in Alfonsa et al. (2015). Specifically, it was necessary to assess the level of jitter amongst the neuronal population with reference to the underlying high frequency (75 to 300 Hz) oscillation.

To do so, spike times were calculated through the equivalent spike detection methods as described in section 2.2.1, in the high (300 Hz to 5 kHz), band pass filtered signal. The Hilbert transform of the lower (75 to 300 Hz), band pass filtered signal was calculated, from which the phase of the oscillation at the time of each spike in the multi-unit activity could be calculated, through the real component of the transform.

These data were used in order to assess the level of the aforementioned jitter of spike timings with reference to the intracellular chloride concentration ($[Cl^-]_i$) of the pyramidal cells. In the experiments performed by Alfonsa et al., the optogenetic chloride ion pump, Halorhodopsin (eNpHR), had been expressed in the pyramidal cells, and so optogenetic illumination was used to putatively increase the $[Cl^-]_i$ within the pyramidal population. The multi-unit activity was then assessed, as described, after a brief electrical stimulation to the white matter, with reference to whether the slice had previously been illuminated, or not (control).

Histograms of the oscillatory phase at spike times were calculated. I helped derive a "half-width index" in order to assess the level of jitter and out-of-phase firing in the population. This was calculated by fitting a Gaussian curve to the results during control (no illumination), from which the half-width of the population response could be calculated, as seen in Figure A.1 (C). This half-width was then applied to the post-illumination results also, within slice, from which the "half-width index" was calculated as follows:

$$Half - width\ index = \frac{\sum(spikes\ outside\ HW)}{\sum(spikes\ within\ HW)}$$

where HW is the half-width measured on the baseline period. To clarify, the HW was calculated only on the baseline epoch, to be applied on both the baseline and the post-illumination data, so as to avoid the effect of a broadening half width due to any potential jitter in the increase $[Cl^-]_i$ situation.

Notice that the frequency band for the spike detection and the frequency band for the underlying oscillation meet, at 300 Hz. Therefore, the resultant alterations to phase may be purely a result of the lower frequency components in the spikes, due to spectral leak. In order to counter this, I helped design two control analyses.

In the first, the spikes were artificially removed from the signal through a simple "amputation" back to the voltages either side of the detected spike. In the second method, the resultant phase was instead calculated at a slight time shift, prior to the occurrence of the spike to nullify the impact of the deflection, and avoid potential postsynaptic effects in the signal. In the time-shift method, the magnitude of the alteration in time was steadily increased until a time when both the amputated and un-amputated signals were equivalent. This indicated a time point at which the underlying oscillation was unaffected by the presence of spikes, with the time ranging between 1.3 and 1.8 ms prior to the spike. Examples of these methods can be seen in Figure A.2. These methods showed that, whilst there was a small effect on the lower frequencies as a result of the spikes in the signal, the results remained significant beyond this effect.

Alfonsa, H., Merricks, E.M., Codadu, N.K., Cunningham, M.O., Deisseroth, K., Racca, C. & Trevelyan, A.J. (2015) The Contribution of Raised Intraneuronal Chloride to Epileptic Network Activity. *Journal of Neuroscience*. 35 (20), 7715–7726.

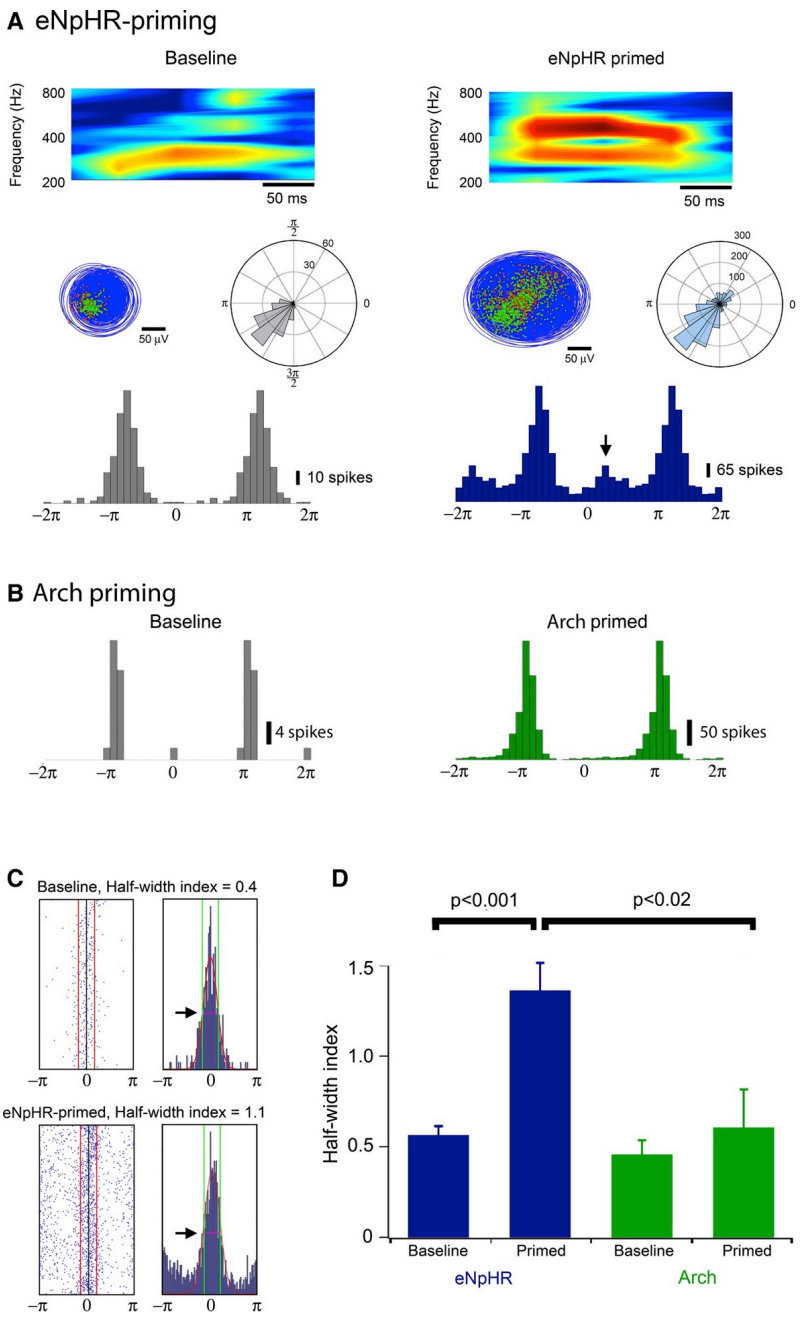


**Figure A.1. Neuronal chloride loading triggers out-of-phase firing during spontaneous bursts of activity**
**A**. Spectrograms of extracellular recordings of spontaneous bursts of activity in baseline (left) and eNpHR-primed (chloride-loaded) tissue. Note the prominent "double-frequency" signal in the chloride-loaded tissue. Detected spikes [dots, red (early) to green transform of the dominant oscillation (75–300 Hz; blue circular trace), and the rose plots rep- esent the numbers of APs occurring at different phases of the oscillation. Note the out-of-phase spiking in the eNpHR-primed dataset, also apparent as a second minor peak (arrowed) in the conventional histogram (duplicated data beyond -π and π). **B**. Periods of hyperpolarization using Arch (Arch-priming) causes a rebound increase in spiking (contrast the calibration bars for the histograms) but no change in the phase distribution of spikes. ***C***, Half-width index measured the ratio between the number of out-of-phase spikes to the in-phase spikes, which were taken to be the spikes within bounds (red lines) set by the half-width of a Gaussian fit to the main spike peak in the baseline histograms. Left column shows the raster plots for baseline and Cl-loaded tissue, and the right column shows the same data plotted as histograms of the spike times. ***D***, Pooled data for half-width indices, showing significantly higher values for Cl-loaded slices (eNpHR-primed) compared with baseline ($n$ = 8, $p < 0.001$, $t$ test) and for Arch-primed tissue ($n$ = 6, $p < 0.02$, $t$ test). Halo, Halorhodopsin. **Figure reproduced from Alfonsa et al. (2015).**

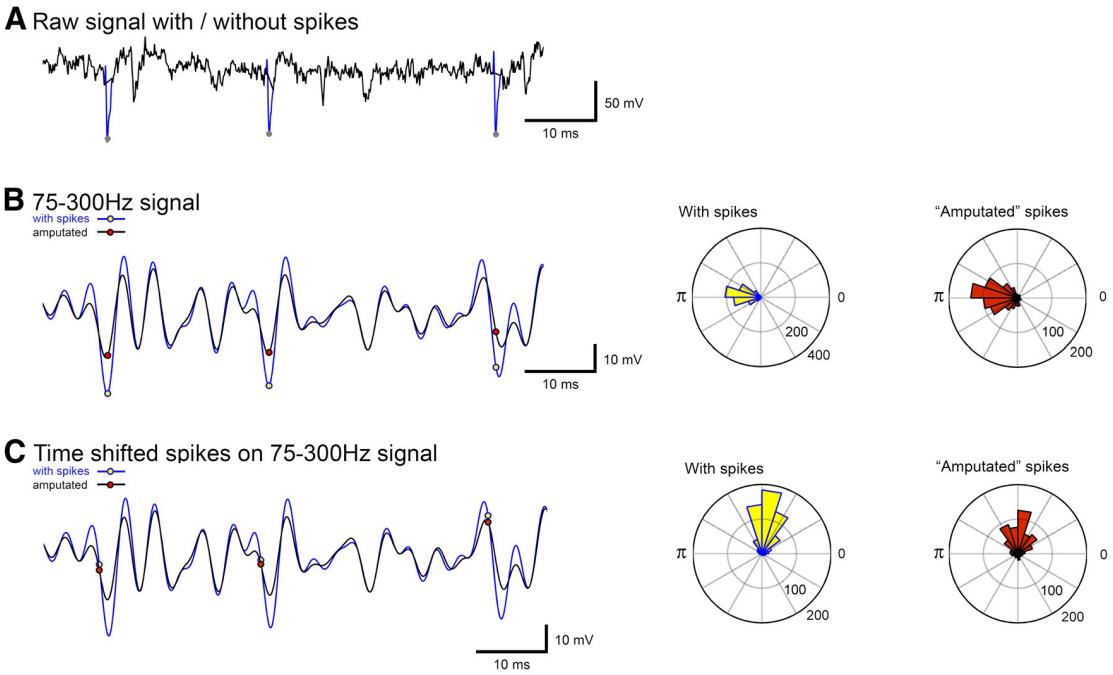


**Figure A.2. Control analyses to examine the effect of spectral leak.**
**A**. A raw data trace, showing the spikes (blue) and the amputated version (black). The bottom traces show the 75–300 Hz bandpass filtered traces. **B**. 75–300 Hz bandpass filtered traces of the raw (blue) and the amputated traces. Red dots indicate the spike times. Note the reduced amplitude peaks for the amputated spikes but that these still occur at the same phase of the oscillation, as evidence by the same skewed orientation of the pooled data shown in the spike-phase plots (right columns). **C.** The same filtered traces but showing the new locations of APs for the time-shifted analysis. The spike-phase plots showed a shifted phase, reflecting the time shift, and a broader main peak, but importantly, they were still heavily skewed for both the raw and amputated data. **Figure reproduced from Alfonsa et al. (2015).**

# References


Aad, G., Abbott, B., Abdallah, J., Abdinov, O., Aben, R., Abolins, M., AbouZeid, O.S., Abramowicz, H., Abreu, H., Abreu, R., Abulaiti, Y., Acharya, B.S., Adamczyk, L., Adams, D.L., Adelman, J., Adomeit, S., Adye, T., Affolder, A.A., Agatonovic-Jovin, T., et al. (2015) Study of the spin and parity of the Higgs boson in diboson decays with the ATLAS detector. *The European Physical Journal. C, Particles and fields*. 75 (10), 476.

Adamos, D.A., Kosmidis, E.K. & Theophilidis, G. (2008) Performance evaluation of PCA-based spike sorting algorithms. *Computer Methods and Programs in Biomedicine*. 91 (3), 232–244.

Adesnik, H. & Scanziani, M. (2010) Lateral competition for cortical space by layer-specific horizontal circuits. *Nature*. 464 (7292), 1155–1160.

Ainsworth, M., Lee, S., Cunningham, M.O., Roopun, A.K., Traub, R.D., Kopell, N.J. & Whittington, M.A. (2011) Dual Gamma Rhythm Generators Control Interlaminar Synchrony in Auditory Cortex. *Journal of Neuroscience*. 31 (47), 17040–17051.

Albowitz, B. & Kuhnt, U. (1995) Epileptiform activity in the guinea-pig neocortical slice spreads preferentially along supragranular layers--recordings with voltage-sensitive dyes. *European Journal of Neuroscience*. 7 (6), 1273–1284.

Aldenkamp, A.P., Overweg, J., Gutter, T., Beun, A.M., Diepman, L. & Mulder, O.G. (1996) Effect of epilepsy, seizures and epileptiform EEG discharges on cognitive function. *Acta neurologica Scandinavica*. 93 (4), 253–259.

Alfonsa, H., Merricks, E.M., Codadu, N.K., Cunningham, M.O., Deisseroth, K., Racca, C. & Trevelyan, A.J. (2015) The Contribution of Raised Intraneuronal Chloride to Epileptic Network Activity. *Journal of Neuroscience.* 35 (20), 7715–7726.

Alvarado-Rojas, C., Lehongre, K., Bagdasaryan, J., Bragin, A., Staba, R., Engel, J., Navarro, V. & Le Van Quyen, M. (2013) Single-unit activities during epileptic discharges in the human hippocampal formation. *Frontiers in Computational Neuroscience*. 7140.

Atallah, B.V., Bruns, W., Carandini, M. & Scanziani, M. (2012) Parvalbumin-Expressing Interneurons Linearly Transform Cortical Responses to Visual Stimuli. *Neuron*. 73 (1), 159–170.

Ayala, G.F., Dichter, M., Gumnit, R.J., Matsumoto, H. & Spencer, W.A. (1973) Genesis of epileptic interictal spikes. New knowledge of cortical feedback systems suggests a neurophysiological explanation of brief paroxysms. *Brain Res*. 521–17.

Babb, T.L. & Crandall, P.H. (1976) Epileptogenesis of human limbic neurons in psychomotor epileptics. *Electroencephalogr Clin Neurophysiol*. 40 (3), 225–243.

Babb, T.L., Wilson, C.L. & Isokawa-Akesson, M. (1987) Firing patterns of human limbic neurons during stereoencephalography (SEEG) and clinical temporal lobe seizures. *Electroencephalogr Clin Neurophysiol*. 66 (6), 467–482.

Baker, P.F., Hodgkin, A.L. & Ridgway, E.B. (1971) Depolarization and calcium entry in squid giant axons. *The Journal of Physiology*. 218 (3), 709–755.

Bar-Hillel, A., Spiro, A. & Stark, E. (2006) Spike sorting: Bayesian clustering of non-stationary data. *Journal of Neuroscience Methods*. 157 (2), 303–316.

Barrese, J.C., Rao, N., Paroo, K., Triebwasser, C., Vargas-Irwin, C., Franquemont, L. & Donoghue, J.P. (2013) Failure mode analysis of silicon-based intracortical microelectrode arrays in non-human primates. *Journal of Neural Engineering*. 10 (6), 066014.

Barthó, P., Hirase, H., Monconduit, L., Zugaro, M., Harris, K.D. & Buzsàki, G. (2004) Characterization of neocortical principal cells and interneurons by network interactions and extracellular features. *Journal of Neurophysiology*. 92 (1), 600–608.

Belichenko, P.V., Sourander, P., Malmgren, K., Nordborg, C., Essen, von, C., Rydenhag, B., Lindström, S., Hedström, A., Uvebrant, P. & Dahlström, A. (1994) Dendritic morphology in epileptogenic cortex from TRPE patients, revealed by intracellular Lucifer Yellow microinjection and confocal laser scanning microscopy. *Epilepsy Research*. 18 (3), 233–247.

Belluscio, M.A., Mizuseki, K. & Schmidt, R. (2012) Cross-frequency phase–phase coupling between theta and gamma oscillations in the hippocampus. *Journal of Neuroscience.* 32 (2), 423–35.

Berkovic, S.F. (2000) SPECT: neurobiology of periictal blood flow alterations. *Advances in Neurology*. 83, 33–39.

Besle, J., Schevon, C.A., Mehta, A.D., Lakatos, P., Goodman, R.R., McKhann, G.M., Emerson, R.G. & Schroeder, C.E. (2011) Tuning of the human neocortex to the temporal dynamics of attended events. *Journal of Neuroscience*. 31 (9), 3176–3185.

Bhatt, D.H., Zhang, S. & Gan, W.-B. (2009) Dendritic Spine Dynamics. *Annual Review of Physiology*. 71 (1), 261–282.

Bialer, M. & White, H.S. (2010) Key factors in the discovery and development of new antiepileptic drugs. *Nature Reviews Drug Discovery*. 9 (1), 68–82.

Bikson, M., Hahn, P.J., Fox, J.E. & Jefferys, J.G.R. (2003) Depolarization block of neurons during maintenance of electrographic seizures. *J Neurophysiol*. 90 (4), 2402–2408.

Blanco, J.A., Stead, M., Krieger, A., Stacey, W., Maus, D., Marsh, E., Viventi, J., Lee, K.H., Marsh, R., Litt, B. & Worrell, G.A. (2011) Data mining neocortical high-frequency oscillations in epilepsy and controls. *Brain*. 134 (10), 2948–2959.

Blümcke, I., Zuschratter, W., Schewe, J.C., Suter, B., Lie, A.A., Riederer, B.M., Meyer, B., Schramm, J., Elger, C.E. & Wiestler, O.D. (1999) Cellular pathology of hilar neurons in Ammon's horn sclerosis. *The Journal of Comparative Neurology*. 414 (4), 437–453.

Bokil, H., Andrews, P., Kulkarni, J.E., Mehta, S. & Mitra, P.P. (2010) Chronux: A platform for analyzing neural signals. *Journal of Neuroscience Methods*. 192 (1), 146–151.

Bothwell, S., Meredith, G.E., Phillips, J., Staunton, H., Doherty, C., Grigorenko, E., Glazier, S., Deadwyler, S.A., O'Donovan, C.A. & Farrell, M. (2001) Neuronal hypertrophy in the neocortex of patients with temporal lobe epilepsy. *Journal of Neuroscience*. 21 (13), 4789–4800.

Bower, M.R. & Buckmaster, P.S. (2008) Changes in granule cell firing rates precede locally recorded spontaneous seizures by minutes in an animal model of temporal lobe epilepsy. *J Neurophysiol*. 99 (5), 2431–2442.

Bower, M.R., Stead, M., Bower, R.S., Kucewicz, M.T., Sulc, V., Cimbalnik, J., Brinkmann, B.H., Vasoli, V.M., St Louis, E.K., Meyer, F.B., Marsh, W.R. & Worrell, G.A. (2015) Evidence for consolidation of neuronal assemblies after seizures in humans. *Journal of Neuroscience*. 35 (3), 999–1010.

Bower, M.R., Stead, M., Meyer, F.B., Marsh, W.R. & Worrell, G.A. (2012) Spatiotemporal neuronal correlates of seizure generation in focal epilepsy. *Epilepsia*. 53 (5), 807–816.

Bragin, A., Penttonen, M. & Buzsàki, G. (1997) Termination of epileptic afterdischarge in the hippocampus. *Journal of Neuroscience*. 17 (7), 2567–2579.

Bravais, L.F. (1827) *Recherches sur les symptômes et le traitement de l'épilepsie hémiplégique*. Paris.

Buzsàki, G. (2006) *Rhythms of the Brain*. Oxford University Press.

Buzsàki, G. & Draguhn, A. (2004) Neuronal oscillations in cortical networks. *Science*. 304 (5679), 1926–1929.

Buzsàki, G. & Watson, B.O. (2012) Brain rhythms and neural syntax: implications for efficient coding of cognitive content and neuropsychiatric disease. *Dialogues in Clinical Neuroscience*. 14 (4), 345–367.

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

Buzsàki, G., Logothetis, N. & Singer, W. (2013) Scaling Brain Size, Keeping Timing: Evolutionary Preservation of Brain Rhythms. *Neuron*. 80 (3), 751–764.

Cammarota, M., Losi, G., Chiavegato, A., Zonta, M. & Carmignoto, G. (2013) Fast spiking interneuron control of seizure propagation in a cortical slice model of focal epilepsy. *The Journal of Physiology*. 591 (Pt 4), 807–822.

Campe, von, G., Spencer, D.D. & de Lanerolle, N.C. (1997) Morphology of dentate granule cells in the human epileptogenic hippocampus. *Hippocampus*. 7 (5), 472–488.

Casanova, J.-L., Conley, M.E., Seligman, S.J., Abel, L. & Notarangelo, L.D. (2014) Guidelines for genetic studies in single patients: lessons from primary immunodeficiencies. *The Journal of Experimental Medicine*. 211 (11), 2137–2149.

Casanova, J.R., Nishimura, M., Le, J., Lam, T.T. & Swann, J.W. (2013) Rapid hippocampal network adaptation to recurring synchronous activity - a role for calcineurin. *European Journal of Neuroscience*. 38 (8), 3115–3127.

Caspers, H. & Speckmann, E.J. (1972) Cerebral pO 2 , pCO 2 and pH: changes during convulsive activity and their significance for spontaneous arrest of seizures. *Epilepsia*. 13 (5), 699–725.

Chagnac-Amitai, Y. & Connors, B.W. (1989) Horizontal spread of synchronized activity in neocortex and its control by GABA-mediated inhibition. *Journal of Neurophysiology*. 61 (4), 747–758.

Clasadonte, J., Dong, J., Hines, D.J. & Haydon, P.G. (2013) Astrocyte control of synaptic NMDA receptors contributes to the progressive development of temporal lobe epilepsy. *Proceedings of the National Academy of Sciences of the United States of America*. 110 (43), 17540–17545.

Cockerell, O.C., Moriarty, J., Trimble, M., Sander, J.W. & Shorvon, S.D. (1996) Acute psychological disorders in patients with epilepsy: a nation-wide study. *Epilepsy Research*. 25 (2), 119–131.

Cohen, I., Navarro, V., Clemenceau, S., Baulac, M. & Miles, R. (2002) On the origin of interictal activity in human temporal lobe epilepsy in vitro. *Science*. 298 (5597), 1418–1421.

Connors, B.W. & Gutnick, M.J. (1990) Intrinsic firing patterns of diverse neocortical neurons. *Trends in Neurosciences*. 13 (3), 99–104.

Craig, M.T. & McBain, C.J. (2015) Fast Gamma Oscillations Are Generated Intrinsically in CA1 without the Involvement of Fast-Spiking Basket Cells. *Journal of Neuroscience*. 35 (8), 3616–3624.

Csercsa, R., Dombovari, B., Fabo, D., Wittner, L., Eross, L., Entz, L., Solyom, A., Rasonyi, G., Szucs, A., Kelemen, A., Jakus, R., Juhos, V., Grand, L., Magony, A., Halasz, P., Freund, T.F., Magloczky, Z., Cash, S.S., Papp, L., et al. (2010) Laminar analysis of slow wave activity in humans. *Brain*. 133 (9), 2814–2829.

Csicsvari, J., Hirase, H., Czurkó, A. & Buzsàki, G. (1998) Reliability and state dependence of pyramidal cell-interneuron synapses in the hippocampus: an ensemble approach in the behaving rat. *Neuron*. 21 (1), 179–189.

Csicsvari, J., Hirase, H., Czurkó, A., Mamiya, A. & Buzsàki, G. (1999) Oscillatory coupling of hippocampal pyramidal cells and interneurons in the behaving Rat. *Journal of Neuroscience*. 19 (1), 274–287.

Cunningham, M.O., Roopun, A., Schofield, I.S., Whittaker, R.G., Duncan, R., Russell, A., Jenkins, A., Nicholson, C., Whittington, M.A. & Traub, R.D. (2012) Glissandi: transient fast electrocorticographic oscillations of steadily increasing frequency, explained by temporally increasing gap junction conductance. *Epilepsia*. 53 (7), 1205–1214.

Destexhe, A., Contreras, D. & Steriade, M. (1999) Spatiotemporal analysis of local field potentials and unit discharges in cat cerebral cortex during natural wake and sleep states. *Journal of Neuroscience*. 19 (11), 4595–4608.

Devergnas, A., Piallat, B., Prabhu, S., Torres, N., Louis Benabid, A., David, O. & Chabardes, S. (2012) The subcortical hidden side of focal motor seizures: evidence from micro-recordings and local field potentials. *Brain*. 135 (7), 2263–2276.

Dichter, M. & Spencer, W.A. (1969a) Penicillin-induced interictal discharges from the cat hippocampus. I. Characteristics and topographical features. *J Neurophysiol*. 32 (5), 649–662.

Dichter, M. & Spencer, W.A. (1969b) Penicillin-induced interictal discharges from the cat hippocampus. II. Mechanisms underlying origin and restriction. *J Neurophysiol*. 32 (5), 663–687.

Donoghue, J.P., Nurmikko, A., Black, M. & Hochberg, L.R. (2007) Assistive technology and robotic control using motor cortex ensemble-based neural interface systems in humans with tetraplegia. *The Journal of Physiology*. 579 (3), 603–611.

Donoho, D.L. & Johnstone, I.M. (1994) Ideal spatial adaptation by wavelet shrinkage. *Biometrika*. 81 (3), 425–455.

Douglas, R., Markram, H. & Martin, K. (2004) 'Neocortex', in Gordon M Shepherd (ed.) *The Synaptic Organisation of the Brain*. 5th edition. Oxford University Press. pp. 449–558.

Dzhala, V.I., Kuchibhotla, K.V., Glykys, J.C., Kahle, K.T., Swiercz, W.B., Feng, G., Kuner, T., Augustine, G.J., Bacskai, B.J. & Staley, K.J. (2010) Progressive NKCC1-dependent neuronal chloride accumulation during neonatal seizures. *Journal of Neuroscience*. 30 (35), 11745–11761.

Ehrens, D., Sritharan, D. & Sarma, S.V. (2015) Closed-loop control of a fragile network: application to seizure-like dynamics of an epilepsy model. *Frontiers in Neuroscience*. 958.

Elger, C.E., Helmstaedter, C. & Kurthen, M. (2004) Chronic epilepsy and cognition. *The Lancet. Neurology*. 3 (11), 663–672.

Ellender, T.J., Raimondo, J.V., Irkle, A., Lamsa, K.P. & Akerman, C.J. (2014) Excitatory Effects of Parvalbumin-Expressing Interneurons Maintain Hippocampal Epileptiform Activity via Synchronous Afterdischarges. *Journal of Neuroscience*. 34, 15208-15222.

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

Epilepsy Foundation of America (1999) *Epilepsy, a Report to the Nation*. Landover, MD.

Erisir, A., Lau, D., Rudy, B. & Leonard, C.S. (1999) Function of specific K(+) channels in sustained high-frequency firing of fast-spiking neocortical interneurons. *J Neurophysiol*. 82 (5), 2476–2489.

Farrell, M.A., Blümcke, I., Khanlou, N. & Vinters, H.V. (2008) 'General Neuropathology of Epilepsy', in Jerome Engel, Jr & Timothy A Pedley (ed.) *Epilepsy: A Comprehensive Textbook*. 2nd edition. Philadelphia: Lippincott-Raven. pp. 103–120.

Fauser, S., Wuwer, Y., Gierschner, C. & Schulze-Bonhage, A. (2004) The localizing and lateralizing value of ictal/postictal coughing in patients with focal epilepsies. *Seizure*. 13 (6), 403–410.

Fazel, S., Wolf, A., Långström, N., Newton, C.R. & Lichtenstein, P. (2013) Premature mortality in epilepsy and the role of psychiatric comorbidity: a total population study. *Lancet*. 382 (9905), 1646–1654.

Fee, M.S., Mitra, P.P. & Kleinfeld, D. (1996) Automatic sorting of multiple unit neuronal signals in the presence of anisotropic and non-Gaussian variability. *Journal of Neuroscience Methods*. 69 (2), 175–188.

Fisher, R., Salanova, V., Witt, T., Worth, R., Henry, T., Gross, R., Oommen, K., Osorio, I., Nazzaro, J., Labar, D., Kaplitt, M., Sperling, M., Sandok, E., Neal, J., Handforth, A., Stern, J., DeSalles, A., Chung, S., Shetter, A., et al. (2010) Electrical stimulation of the anterior nucleus of thalamus for treatment of refractory epilepsy. *Epilepsia*. 51 (5), 899–908.

Fisher, R.A. (1936) The Use of Multiple Measurements in Taxonomic Problems. *Annals of Eugenics*. 7 (2), 179–188.

Fisher, R.S. & Prince, D.A. (1977) Spike-wave rhythms in cat cortex induced by parenteral penicillin. II. Cellular features. *Electroencephalography and Clinical Neurophysiology*. 42 (5), 625–639.

Fisher, R.S. & Schachter, S.C. (2000) The Postictal State: A Neglected Entity in the Management of Epilepsy. *Epilepsy & Behavior*. 1 (1), 52–59.

Fisher, R.S., Acevedo, C., Arzimanoglou, A., Bogacz, A., Cross, J.H., Elger, C.E., Engel, J., Forsgren, L., French, J.A., Glynn, M., Hesdorffer, D.C., Lee, B.I., Mathern, G.W., Moshé, S.L., Perucca, E., Scheffer, I.E., Tomson, T., Watanabe, M. & Wiebe, S. (2014) ILAE official report: a practical clinical definition of epilepsy. *Epilepsia*. 55 (4), 475–482.

Freund, T.F. & Buzsàki, G. (1996) Interneurons of the hippocampus. *Hippocampus*. 6 (4), 347–470.

Fried, I., Wilson, C.L., Maidment, N.T., Engel, J., Behnke, E., Fields, T.A., MacDonald, K.A., Morrow, J.W. & Ackerson, L. (1999) Cerebral microdialysis combined with single-neuron and electroencephalographic recording in neurosurgical patients. Technical note. *Journal of Neurosurgery*. 91 (4), 697–705.

Fujiwara-Tsukamoto, Y., Isomura, Y. & Takada, M. (2006) Comparable GABAergic Mechanisms of Hippocampal Seizurelike Activity in Posttetanic and Low-Mg2+ Conditions. *Journal of Neurophysiology.* 95, 2013-2019.

Gadhoumi, K., Lina, J.-M., Mormann, F. & Gotman, J. (2015) Seizure prediction for therapeutic devices: A review. *Journal of Neuroscience Methods*. doi: 10.1016/j.jneumeth.2015.06.010.

Gallmetzer, P., Leutmezer, F., Serles, W., Assem-Hilger, E., Spatt, J. & Baumgartner, C. (2004) Postictal paresis in focal epilepsies--incidence, duration, and causes: a video-EEG monitoring study. *Neurology*. 62 (12), 2160–2164.

Gilja, V., Linderman, M.D., Santhanam, G., Afshar, A., Ryu, S., Meng, T.H. & Shenoy, K.V. (2006) Multiday Electrophysiological Recordings from Freely Behaving Primates. *Conf Proc IEEE Eng Med Biol Soc.* 1, 5643–5646.

Gold, C., Gold, C., Henze, D.A., Henze, D.A., Koch, C., Buzsàki, G. & Buzsáki, G. (2006) On the origin of the extracellular action potential waveform: A modeling study. *J Neurophysiol*. 95 (5), 3113–3128.

Grone, B.P. & Baraban, S.C. (2015) Animal models in epilepsy research: legacies and new directions. *Nature Neuroscience*. 18 (3), 339–343.

Gummadavelli, A., Motelow, J.E., Smith, N., Zhan, Q., Schiff, N.D. & Blumenfeld, H. (2014) Thalamic stimulation to improve level of consciousness after seizures: Evaluation of electrophysiology and behavior. *Epilepsia*. 56 (1), 114-24.

Gumnit, R.J. & Takahashi, T. (1965) Changes in direct current activity during experimental focal seizures. *Electroencephalogr Clin Neurophysiol*. 19 (1), 63–74.

Guo, D., Arnspiger, S., Rensing, N.R. & Wong, M. (2012) Brief seizures cause dendritic injury. *Neurobiology of Disease*. 45 (1), 348–355.

Hammers, A., Asselin, M.-C., Hinz, R., Kitchen, I., Brooks, D.J., Duncan, J.S. & Koepp, M.J. (2007) Upregulation of opioid receptor binding following spontaneous epileptic seizures. *Brain*. 130 (Pt 4), 1009–1016.

Harris, K.D., Henze, D.A., Csicsvari, J., Hirase, H. & Buzsàki, G. (2000) Accuracy of tetrode spike separation as determined by simultaneous intracellular and extracellular measurements. *J Neurophysiol*. 84 (1), 401–414.

Hauser, W.A., Annegers, J.F. & Kurland, L.T. (1993) Incidence of epilepsy and unprovoked seizures in Rochester, Minnesota: 1935-1984. *Epilepsia*. 34 (3), 453–468.

Helfrich, R.F., Knepper, H., Nolte, G., Strüber, D., Rach, S., Herrmann, C.S., Schneider, T.R. & Engel, A.K. (2014) Selective Modulation of Interhemispheric Functional Connectivity by HD-tACS Shapes Perception. *PLoS Biology*. 12 (12), e1002031.

Helmstaedter, C., Elger, C.E. & Lendt, M. (1994) Postictal courses of cognitive deficits in focal epilepsies. *Epilepsia*. 35 (5), 1073–1078.

Hendry, S.H., Schwark, H.D., Jones, E.G. & Yan, J. (1987) Numbers and proportions of GABA-immunoreactive neurons in different areas of monkey cerebral cortex. *Journal of Neuroscience*. 7 (5), 1503–1519.

Henze, D.A., Borhegyi, Z., Csicsvari, J., Mamiya, A., Harris, K.D. & Buzsàki, G. (2000) Intracellular features predicted by extracellular recordings in the hippocampus in vivo. *J Neurophysiol*. 84 (1), 390–400.

Herbst, J.A., Gammeter, S., Ferrero, D. & Hahnloser, R.H.R. (2008) Spike sorting with hidden Markov models. *Journal of Neuroscience Methods*. 174 (1), 126–134.

Hesdorffer, D.C., Logroscino, G., Benn, E.K.T., Katri, N., Cascino, G. & Hauser, W.A. (2011) Estimating risk for developing epilepsy: a population-based study in Rochester, Minnesota. *Neurology*. 76 (1), 23–27.

Hill, D.N., Mehta, S.B. & Kleinfeld, D. (2011) Quality metrics to accompany spike sorting of extracellular signals. *Journal of Neuroscience*. 31 (24), 8699–8705.

Hochberg, L.R., Serruya, M.D., Friehs, G.M., Mukand, J.A., Saleh, M., Caplan, A.H., Branner, A., Chen, D., Penn, R.D. & Donoghue, J.P. (2006) Neuronal ensemble control of prosthetic devices by a human with tetraplegia. *Nature*. 442 (7099), 164–171.

Hoffmann, J.M., Elger, C.E. & Kleefuss-Lie, A.A. (2009) The localizing value of hypersalivation and postictal coughing in temporal lobe epilepsy. *Epilepsy Research*. 87 (2-3), 144–147.

Holmes, M.D., Brown, M. & Tucker, D.M. (2004) Are 'generalized' seizures truly generalized? Evidence of localized mesial frontal and frontopolar discharges in absence. *Epilepsia*. 45 (12), 1568–1579.

Holst, A.G., Winkel, B.G., Risgaard, B., Nielsen, J.B., Rasmussen, P.V., Haunsø, S., Sabers, A., Uldall, P. & Tfelt-Hansen, J. (2013) Epilepsy and risk of death and sudden unexpected death in the young: a nationwide study. *Epilepsia*. 54 (9), 1613–1620.

Hubel, D.H. & Wiesel, T.N. (1977) Ferrier lecture. Functional architecture of macaque monkey visual cortex. *Proc R Soc Lond B Biol Sci*. 198 (1130), 1–59.

Huberfeld, G., Wittner, L., Clemenceau, S., Baulac, M., Kaila, K., Miles, R. & Rivera, C. (2007) Perturbed chloride homeostasis and GABAergic signaling in human temporal lobe epilepsy. *Journal of Neuroscience*. 27 (37), 9866–9873.

Ikeda, A., Taki, W., Kunieda, T., Terada, K., Mikuni, N., Nagamine, T., Yazawa, S., Ohara, S., Hori, T., Kaji, R., Kimura, J. & Shibasaki, H. (1999) Focal ictal direct current shifts in human epilepsy as studied by subdural and scalp recording. *Brain*. 122 ( Pt 5)827–838.

Ikeda, A., Terada, K., Mikuni, N., Burgess, R.C., Comair, Y., Taki, W., Hamano, T., Kimura, J., Lüders, H.O. & Shibasaki, H. (1996) Subdural recording of ictal DC shifts in neocortical seizures in humans. *Epilepsia*. 37 (7), 662–674.

Isaacson, J.S. & Scanziani, M. (2011) How inhibition shapes cortical activity. *Neuron*. 72 (2), 231–243.

Ishijima, B. (1972) Unitary analysis of epileptic activity in acute and chronic foci and related cortex of cat and monkey. *Epilepsia*. 13 (4), 561–581.

Isokawa, M. (1997) Preservation of dendrites with the presence of reorganized mossy fiber collaterals in hippocampal dentate granule cells in patients with temporal lobe epilepsy. *Brain Research*. 744 (2), 339–343.

Jack, J., Noble, D. & Tsien, R.W. (1975) *Electric current flow in excitable cells*. Oxford University Press.

Jackson, J.H. (1879) Lectures on the Diagnosis of Epilepsy. *British Medical Journal*. 1 (943), 109–112.

Jackson, J.H. (1875) *On temporary mental disorders after epileptic paroxysms*. West Riding Lunatic Assylum Medical Reports.

Jacobs, J., Kahana, M.J., Ekstrom, A.D. & Fried, I. (2007) Brain oscillations control timing of single-neuron activity in humans. *Journal of Neuroscience*. 27 (14), 3839–3844.

Jaffe, D.B., Johnston, D., Lasser-Ross, N., Lisman, J.E., Miyakawa, H. & Ross, W.N. (1992) The spread of Na+ spikes determines the pattern of dendritic Ca2+ entry into hippocampal neurons. *Nature*. 357 (6375), 244–246.

Jan, M.M.S., Sadler, M. & Rahey, S.R. (2001) Lateralized Postictal EEG Delta Predicts the Side of Seizure Surgery in Temporal Lobe Epilepsy. *Epilepsia*. 42 (3), 1–4.

Janszky, J., Fogarasi, A., Toth, V., Magalova, V., Gyimesi, C., Kovacs, N., Schulz, R. & Ebner, A. (2007) Peri-ictal vegetative symptoms in temporal lobe epilepsy. *Epilepsy & Behavior*. 11 (1), 125–129.

Jones, R.S.G., da Silva, A.B., Whittaker, R.G., Woodhall, G.L. & Cunningham, M.O. (2015) Human brain slices for epilepsy research: Pitfalls, solutions and future challenges. *Journal of Neuroscience Methods*. http://dx.doi.org/10.1016/j.jneumeth.2015.09.021

Kahle, K.T., Staley, K.J., Nahed, B.V., Gamba, G., Hebert, S.C., Lifton, R.P. & Mount, D.B. (2008) Roles of the cation-chloride cotransporters in neurological disease. *Nature Clinical Practice Neurology*. 4 (9), 490–503.

Kaibara, M. & Blume, W.T. (1988) The postictal electroencephalogram. *Electroencephalography and Clinical Neurophysiology*. 70 (2), 99–104.

Kaila, K., Price, T.J., Payne, J.A., Puskarjov, M. & Voipio, J. (2014) Cation-chloride cotransporters in neuronal development, plasticity and disease. *Nature Reviews Neuroscience*. 15 (10), 637–654.

Kandel, E.R. & Spencer, W.A. (1961a) Electrophysiology of hippocampal neurons. II. After-potentials and repetitive firing. *J Neurophysiol*. 24243–259.

Kandel, E.R. & Spencer, W.A. (1961b) The pyramidal cell during hippocampal seizure. *Epilepsia*. 2 (1), 63–69.

Kanemoto, K., Kawasaki, J. & Mori, E. (1999) Violence and epilepsy: a close relation between violence and postictal psychosis. *Epilepsia*. 40 (1), 107–109.

Karnani, M.M., Agetsuma, M. & Yuste, R. (2014) A blanket of inhibition: functional inferences from dense inhibitory connectivity. *Current Opinion in Neurobiology*. 2696–102.

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

Keller, C.J., Truccolo, W., Gale, J.T., Eskandar, E., Thesen, T., Carlson, C., Devinsky, O., Kuzniecky, R., Doyle, W.K., Madsen, J.R., Schomer, D.L., Mehta, A.D., Brown, E.N., Hochberg, L.R., Ulbert, I., Halgren, E. & Cash, S.S. (2010) Heterogeneous neuronal firing patterns during interictal epileptiform discharges in the human cortex. *Brain*. 133 (6), 1668–1681.

Khalilov, I., Holmes, G.L. & Ben-Ari, Y. (2003) In vitro formation of a secondary epileptogenic mirror focus by interhippocampal propagation of seizures. *Nature Neuroscience*. 6 (10), 1079–1085.

Khan, O.I., Zhao, Q., Miller, F. & Holmes, G.L. (2010) Interictal spikes in developing rats cause long-standing cognitive deficits. *Neurobiology of Disease*. 39 (3), 362–371.

Kim, W., Miller, J.W., Ojemann, J.G. & Miller, K.J. (2009) Ictal localization by invasive recording of infraslow activity with DC-coupled amplifiers. *Journal of Clinical Neurophysiology*. 26 (3), 135–144.

Klausberger, T., Roberts, J.D.B. & Somogyi, P. (2002) Cell type- and input-specific differences in the number and subtypes of synaptic GABA(A) receptors in the hippocampus. *Journal of Neuroscience*. 22 (7), 2513–2521.

Kleen, J.K., Scott, R.C., Holmes, G.L. & Lenck-Santini, P.P. (2010) Hippocampal interictal spikes disrupt cognition in rats. *Annals of Neurology*. 67 (2), 250–257.

Kramer, M.A., Truccolo, W., Eden, U.T., Lepage, K.Q., Hochberg, L.R., Eskandar, E.N., Madsen, J.R., Lee, J.W., Maheshwari, A., Halgren, E., Chu, C.J. & Cash, S.S. (2012) Human seizures self-terminate across spatial scales via a critical transition. *Proceedings of the National Academy of Sciences of the United States of America*. 109 (51), 21116–21121.

Krauss, G. & Theodore, W.H. (2010) Treatment strategies in the postictal state. *Epilepsy Behav*. 19 (2), 188–190.

Krook-Magnuson, E., Armstrong, C., Oijala, M. & Soltesz, I. (2013) On-demand optogenetic control of spontaneous seizures in temporal lobe epilepsy. *Nature Communications*. 41376–8.

Kros, L., Eelkman Rooda, O.H.J., Spanke, J.K., Alva, P., van Dongen, M.N., Karapatis, A., Tolner, E.A., Strydis, C., Davey, N., Winkelman, B.H.J., Negrello, M., Serdijn, W.A., Steuber, V., van den Maagdenberg, A.M.J.M., De Zeeuw, C.I. & Hoebeek, F.E. (2015) Cerebellar output controls generalized spike-and-wave discharge occurrence. *Annals of Neurology*. 77 (6), 1027–1049.

Kullmann, D.M., Schorge, S., Walker, M.C. & Wykes, R.C. (2014) Gene therapy in epilepsy—is it time for clinical trials? *Nature Reviews Neurology.* 10 (5), 300–304.

Lakatos, P., Karmos, G., Mehta, A.D., Ulbert, I. & Schroeder, C.E. (2008) Entrainment of neuronal oscillations as a mechanism of attentional selection. *Science*. 320 (5872), 110–113.

Laxer, K.D., Trinka, E., Hirsch, L.J., Cendes, F., Langfitt, J., Delanty, N., Resnick, T. & Benbadis, S.R. (2014) The consequences of refractory epilepsy and its treatment. *Epilepsy Behav*. 3759–70.

Le Van Quyen, M., Staba, R., Bragin, A., Dickson, C., Valderrama, M., Fried, I. & Engel, J. (2010) Large-Scale Microelectrode Recordings of High-Frequency Gamma Oscillations in Human Cortex during Sleep. *Journal of Neuroscience*. 30 (23), 7770–7782.

Lee, S., Kruglikov, I., Huang, Z.J., Fishell, G. & Rudy, B. (2013) A disinhibitory circuit mediates motor integration in the somatosensory cortex. *Nature Neuroscience*. 16 (11), 1662–1670.

Lee, S.-H., Kwan, A.C., Zhang, S., Phoumthipphavong, V., Flannery, J.G., Masmanidis, S.C., Taniguchi, H., Huang, Z.J., Zhang, F., Boyden, E.S., Deisseroth, K. & Dan, Y. (2012) Activation of specific interneurons improves V1 feature selectivity and visual perception. *Nature*. 488 (7411), 379–383.

Łęski, S., Lindén, H., Tetzlaff, T., Pettersen, K.H. & Einevoll, G.T. (2013) Frequency dependence of signal power and spatial reach of the local field potential. Olaf Sporns & Olaf Sporns (eds.). *PLoS Computational Biology*. 9 (7), e1003137.

Leutmezer, F. & Baumgartner, C. (2002) Postictal signs of lateralizing and localizing significance. *Epileptic Disorders*. 4 (1), 43–48.

Linderman, M.D., Linderman, M.D., Gilja, V., Gilja, V., Santhanam, G., Santhanam, G., Afshar, A., Afshar, A., Ryu, S., Ryu, S., Meng, T.H., Meng, T.H., Shenoy, K.V. & Shenoy, K.V. (2006) Neural recording stability of chronic electrode arrays in freely behaving primates. *Conference Proceedings IEEE Engineering in Medicine and Biology Society*. 1, 4387–4391.

Löscher, W. (2011) Critical review of current animal models of seizures and epilepsy used in the discovery and development of new antiepileptic drugs. *Seizure*. 20 (5), 359–368.

Löscher, W. & Köhling, R. (2010) Functional, metabolic, and synaptic changes after seizures as potential targets for antiepileptic therapy. *Epilepsy Behav*. 19 (2), 105–113.

Lux, H.D., Heinemann, U. & Dietzel, I. (1986) Ionic changes and alterations in the size of the extracellular space during epileptic activity. *Advances in Neurology*. 44619–639.

MacArthur, D.G., Manolio, T.A., Dimmock, D.P., Rehm, H.L., Shendure, J., Abecasis, G.R., Adams, D.R., Altman, R.B., Antonarakis, S.E., Ashley, E.A., Barrett, J.C., Biesecker, L.G., Conrad, D.F., Cooper, G.M., Cox, N.J., Daly, M.J., Gerstein, M.B., Goldstein, D.B., Hirschhorn, J.N., et al. (2014) Guidelines for investigating causality of sequence variants in human disease. *Nature*. 508 (7497), 469–476.

Martina, M., Schultz, J.H., Ehmke, H., Monyer, H. & Jonas, P. (1998) Functional and molecular differences between voltage-gated K+ channels of fast-spiking interneurons and pyramidal neurons of rat hippocampus. *Journal of Neuroscience*. 18 (20), 8111–8125.

Masse, N.Y., Jarosiewicz, B., Simeral, J.D., Bacher, D., Stavisky, S.D., Cash, S.S., Oakley, E.M., Berhanu, E., Eskandar, E., Friehs, G., Hochberg, L.R. & Donoghue, J.P. (2014) Non-causal spike filtering improves decoding of movement intention for intracortical BCIs. *Journal of Neuroscience Methods*. 23658–67.

Matsumoto, H. & Marsan, C.A. (1964) Cortical cellular phenomena in experimental epilepsy: Ictal manifestations. *Exp Neurol*. 9 (4), 305–326.

Maynard, E.M., Hatsopoulos, N.G., Ojakangas, C.L., Acuna, B.D., Sanes, J.N., Normann, R.A. & Donoghue, J.P. (1999) Neuronal interactions improve cortical population coding of movement direction. *Journal of Neuroscience*. 19 (18), 8083–8093.

McCormick, D.A. & Feeser, H.R. (1990) Functional implications of burst firing and single spike activity in lateral geniculate relay neurons. *Neuroscience*.

McGaugh, J.L. (2000) Memory--a century of consolidation. *Science*. 287 (5451), 248–251.

Megías, M., Emri, Z., Freund, T.F. & Gulyas, A.I. (2001) Total number and distribution of inhibitory and excitatory synapses on hippocampal CA1 pyramidal cells. *Neuroscience*. 102 (3), 527–540.

Merker, B. (2013) Cortical gamma oscillations: the functional key is activation, not cognition. *Neuroscience and Biobehavioral Reviews*. 37 (3), 401–417.

Merricks, E.M., Smith, E.H., McKhann, G.M., Goodman, R.R., Bateman, L.M., Emerson, R.G., Schevon, C.A. & Trevelyan, A.J. (2015) Single unit action potentials in humans and the effect of seizure activity. *Brain*. 138 (10), 2891–906.

Mitchell, J.F., Sundberg, K.A. & Reynolds, J.H. (2007) Differential attention-dependent response modulation across cell classes in macaque visual area V4. *Neuron*. 55 (1), 131–141.

Mizuseki, K., Sirota, A., Pastalkova, E. & Buzsàki, G. (2009) Theta Oscillations Provide Temporal Windows for Local Circuit Computationin the Entorhinal-Hippocampal Loop. *Neuron*. 64 (2), 267–280.

Mormann, F., Andrzejak, R.G., Elger, C.E. & Lehnertz, K. (2007) Seizure prediction: the long and winding road. *Brain*. 130 (2), 314–333.

Moshé, S.L., Perucca, E., Ryvlin, P. & Tomson, T. (2015) Epilepsy: new advances. *Lancet*. 385 (9971), 884–898.

Multani, P., Myers, R.H., Blume, H.W., Schomer, D.L. & Sotrel, A. (1994) Neocortical dendritic pathology in human partial epilepsy: a quantitative Golgi study. *Epilepsia*. 35 (4), 728–736.

Nadler, J.V. & Spencer, D.D. (2014) What is a seizure focus? *Advances in Experimental Medicine and Biology*. 813 (4), 55–62.

Nelson, S.B. & Turrigiano, G.G. (1998) Synaptic depression: a key player in the cortical balancing act. *Nature Neuroscience*. 1 (7), 539–541.

Nicoll, R.A., Malenka, R.C. & Kauer, J.A. (1990) Functional comparison of neurotransmitter receptor subtypes in mammalian central nervous system. *Physiological Reviews*. 70 (2), 513–565.

Niedermeyer, E. & da Silva, F.L. (2005) *Electroencephalography: basic principles, clinical applications, and related fields*. Lippincott Williams & Wilkins.

Nir, Y., Fisch, L., Mukamel, R., Gelbard-Sagiv, H., Arieli, A., Fried, I. & Malach, R. (2007) Coupling between neuronal firing rate, gamma LFP, and BOLD fMRI is related to interneuronal correlations. *Current Biology*. 17 (15), 1275–1285.

Nunez, P.L. (1981) *Electric Fields of the Brain: The Neurophysics of EEG*. Oxford University Press.

O'Keefe, J. & Recce, M.L. (1993) Phase relationship between hippocampal place units and the EEG theta rhythm. *Hippocampus*. 3 (3), 317–330.

Pandarinath, C., Gilja, V., Blabe, C.H. & Nuyujukian, P. (2015) Neural population dynamics in human motor cortex during movements in people with ALS. *eLife*. doi: 10.7554/eLife.07436

Park, Y.S., Hochberg, L.R., Eskandar, E.N., Cash, S.S. & Truccolo, W. (2014) Early detection of human focal seizures based on cortical multiunit activity. *Conference Proceedings IEEE Engineering in Medicine and Biology Society*. 5796–5799.

Paz, J.T., Davidson, T.J., Frechette, E.S., Delord, B., Parada, I., Peng, K., Deisseroth, K. & Huguenard, J.R. (2012) Closed-loop optogenetic control of thalamus as a tool for interrupting seizures after cortical injury. *Nature Neuroscience*. 16 (1), 64–70.

Pearce, J.M. (1994) Robert Bentley Todd (1809-60) and Todd's paralysis. *J Neurol Neurosurg Psychiatry.* 57 (3), 315.

Perge, J.A., Homer, M.L., Malik, W.Q., Cash, S., Eskandar, E., Friehs, G., Donoghue, J.P. & Hochberg, L.R. (2013) Intra-day signal instabilities affect decoding performance in an intracortical neural interface system. *Journal of Neural Engineering*. 10 (3), 036004–26.

Perge, J.A., Zhang, S., Malik, W.Q., Homer, M.L., Cash, S., Friehs, G., Eskandar, E.N., Donoghue, J.P. & Hochberg, L.R. (2014) Reliability of directional information in unsorted spikes and local field potentials recorded in human motor cortex. *Journal of Neural Engineering*. 11 (4), 046007–28.

Peron, S., Chen, T.-W. & Svoboda, K. (2015) Comprehensive imaging of cortical networks. *Current Opinion in Neurobiology*. doi: 10.1016/j.conb.2015.03.016.

Peyrache, A., Dehghani, N., Eskandar, E.N., Madsen, J.R., Anderson, W.S., Donoghue, J.A., Hochberg, L.R., Halgren, E., Cash, S.S. & Destexhe, A. (2012) Spatiotemporal dynamics of neocortical excitation and inhibition during human sleep. *Proceedings of the National Academy of Sciences of the United States of America*. 109 (5), 1731–1736.

Pfeffer, C.K., Xue, M., He, M., Huang, Z.J. & Scanziani, M. (2013) Inhibition of inhibition in visual cortex: the logic of connections between molecularly distinct interneurons. *Nature Neuroscience*. 16 (8), 1068–1076.

Pi, H.-J., Hangya, B., Kvitsiani, D., Sanders, J.I., Huang, Z.J. & Kepecs, A. (2013) Cortical interneurons that specialize in disinhibitory control. *Nature*. 503 (7477), 521–524.

Pinto, D.J., Patrick, S.L., Huang, W.C. & Connors, B.W. (2005) Initiation, propagation, and termination of epileptiform activity in rodent neocortex in vitro involve distinct mechanisms. *Journal of Neuroscience*. 25 (36), 8131–8140.

Pockett, S. & Holmes, M.D. (2009) Intracranial EEG power spectra and phase synchrony during consciousness and unconsciousness. *Consciousness and Cognition*. 18 (4), 1049–1055.

Pouille, F., Watkinson, O., Scanziani, M. & Trevelyan, A.J. (2013) The contribution of synaptic location to inhibitory gain control in pyramidal cells. *Physiological Reports*. 1 (5), e00067–n/a.

Pouzat, C., Mazor, O. & Laurent, G. (2002) Using noise signature to optimize spike-sorting and to assess neuronal classification quality. *Journal of Neuroscience Methods*. 122 (1), 43–57.

Prince, D.A. & Wilder, B.J. (1967) Control mechanisms in cortical epileptogenic foci. 'Surround' inhibition. *Arch Neurol*. 16 (2), 194–202.

Privitera, M., Kohler, C., Cahill, W. & Yeh, H.S. (1996) Postictal language dysfunction in patients with right or bilateral hemispheric language localization. *Epilepsia*. 37 (10), 936–941.

Privitera, M.D., Morris, G.L. & Gilliam, F. (1991) Postictal language assessment and lateralization of complex partial seizures. *Annals of Neurology*. 30 (3), 391–396.

Quian Quiroga, R., Nadasdy, Z. & Ben-Shaul, Y. (2004) Unsupervised Spike Detection and Sorting with Wavelets and Superparamagnetic Clustering. *Neural Computation*. 16 (8), 1661–1687.

Quian Quiroga, R., Reddy, L., Koch, C. & Fried, I. (2007) Decoding Visual Inputs From Multiple Neurons in the Human Temporal Lobe. *J Neurophysiol*. 98 (4), 1997–2007.

Quian Quiroga, R., Reddy, L., Kreiman, G., Koch, C. & Fried, I. (2005) Invariant visual representation by single neurons in the human brain. *Nature*. 435 (7045), 1102–1107.

Rampp, S. & Stefan, H. (2012) Ictal onset baseline shifts and infraslow activity. *Journal of Clinical Neurophysiology*. 29 (4), 291–297.

Ranck, J.B. (1973) Studies on single neurons in dorsal hippocampal formation and septum in unrestrained rats. I. Behavioral correlates and firing repertoires. *Exp Neurol*. 41 (2), 461–531.

Ray, S. & Maunsell, J.H.R. (2011) Different Origins of Gamma Rhythm and High-Gamma Activity in Macaque Visual Cortex. *PLoS Biology*. 9 (4), e1000610.

Rektor, I., Zákopcan, J., Tyrlíková, I., Kuba, R., Brázdil, M., Chrastina, J. & Novák, Z. (2009) Secondary generalization in seizures of temporal lobe origin: Ictal EEG pattern in a stereo-EEG study. *Epilepsy Behav*. 15 (2), 235–239.

Rémi, J. & Noachtar, S. (2010) Clinical features of the postictal state: correlation with seizure variables. *Epilepsy Behav*. 19 (2), 114–117.

Rolak, L.A., Rutecki, P., Ashizawa, T. & Harati, Y. (1992) Clinical features of Todd's post-epileptic paralysis. *J Neurol Neurosurg Psychiatry*. 55 (1), 63–64.

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. & 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.

Rutishauser, U., Schuman, E.M. & Mamelak, A.N. (2006) Online detection and sorting of extracellularly recorded action potentials in human medial temporal lobe recordings, in vivo. *Journal of Neuroscience Methods*. 154 (1-2), 204–224.

Rutishauser, U., Ross, I.B., Mamelak, A.N. & Schuman, E.M. (2010) Human memory strength is predicted by theta-frequency phase-locking of single neurons. *Nature*. 464 (7290), 903–907.

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

Schevon, C.A., Weiss, S.A., McKhann, G., Goodman, R.R., Yuste, R., Emerson, R.G. & Trevelyan, A.J. (2012) Evidence of an inhibitory restraint of seizure activity in humans. *Nature Communications*. 31060–11.

Schindler, K., Leung, H., Lehnertz, K. & Elger, C.E. (2007) How generalised are secondarily 'generalised' tonic clonic seizures? *J Neurol Neurosurg Psychiatry*. 78 (9), 993–996.

Schmidt, E.M., Mutsuga, N. & McIntosh, J.S. (1976) Chronic recording of neurons in epileptogenic foci of monkey during seizures. *Exp Neurol*. 52 (3), 459–466.

Schölvinck, M.L., Saleem, A.B., Benucci, A., Harris, K.D. & Carandini, M. (2015) Cortical state determines global variability and correlations in visual cortex. *Journal of Neuroscience*. 35 (1), 170–178.

Schuchmann, S., Schmitz, D., Rivera, C., Vanhatalo, S., Salmen, B., Mackie, K., Sipilä, S.T., Voipio, J. & Kaila, K. (2006) Experimental febrile seizures are precipitated by a hyperthermia-induced respiratory alkalosis. *Nat Med*. 12 (7), 817–823.

Schwartz, T.H. & Bonhoeffer, T. (2001) In vivo optical mapping of epileptic foci and surround inhibition in ferret cerebral cortex. *Nat Med*. 7 (9), 1063–1067.

Sedley, W. & Cunningham, M.O. (2013) Do cortical gamma oscillations promote or suppress perception? An under-asked question with an over-assumed answer. *Frontiers in Human Neuroscience*. 7, 595.

Seino, M. (2006) Classification criteria of epileptic seizures and syndromes. *Epilepsy Research*. 70 Suppl 1, S27–33.

Shepherd, G.M. (2004) *The Synaptic Organisation of the Brain*. 5th edition. Oxford University Press.

Simeral, J.D., Kim, S.-P., Black, M.J., Donoghue, J.P. & Hochberg, L.R. (2011) Neural control of cursor trajectory and click by a human with tetraplegia 1000 days after implant of an intracortical microelectrode array. *Journal of Neural Engineering*. 8 (2), 025027–25.

Simon, A., Traub, R.D., Vladimirov, N., Jenkins, A., Nicholson, C., Whittaker, R.G., Schofield, I., Clowry, G.J., Cunningham, M.O. & Whittington, M.A. (2013) Gap junction networks can generate both ripple-like and fast ripple-like oscillations. *European Journal of Neuroscience*. 39 (1), 46–60.

Sirota, A., Montgomery, S., Fujisawa, S., Isomura, Y., Zugaro, M. & Buzsàki, G. (2008) Entrainment of neocortical neurons and gamma oscillations by the hippocampal theta rhythm. *Neuron*. 60 (4), 683–697.

Skaggs, W.E., McNaughton, B.L., Wilson, M.A. & Barnes, C.A. (1996) Theta phase precession in hippocampal neuronal populations and the compression of temporal sequences. *Hippocampus*. 6 (2), 149–172.

So, N.K. & Blume, W.T. (2010) The postictal EEG. *Epilepsy & Behavior*. 19 (2), 121–126.

Somjen, G.G. (2004) *Ions in the brain*. New York: Oxford University Press.

Somjen, G.G. & Giacchino, J.L. (1985) Potassium and calcium concentrations in interstitial fluid of hippocampal formation during paroxysmal responses. *Journal of Neurophysiology*. 53, 1098–1108.

Somogyi, P. & Klausberger, T. (2005) Defined types of cortical interneurone structure space and spike timing in the hippocampus. *The Journal of Physiology*. 562 (Pt 1), 9–26.

Stead, M., Bower, M., Brinkmann, B.H., Lee, K., Marsh, W.R., Meyer, F.B., Litt, B., Van Gompel, J. & Worrell, G.A. (2010) Microseizures and the spatiotemporal scales of human partial epilepsy. *Brain*. 133 (9), 2789–2797.

Stegen, M., Young, C.C., Haas, C.A., Zentner, J. & Wolfart, J. (2009) Increased leak conductance in dentate gyrus granule cells of temporal lobe epilepsy patients with Ammon's horn sclerosis. *Epilepsia*. 50 (4), 646–653.

Steriade, M., Amzica, F., Neckelmann, D. & Timofeev, I. (1998) Spike-wave complexes and fast components of cortically generated seizures II. Extra-and intracellular patterns. *Journal of Neurophysiology.* 80 (3), 1456–79.

Steriade, M., McCormick, D.A. & Sejnowski, T.J. (1993) Thalamocortical oscillations in the sleeping and aroused brain. *Science*. 262 (5134), 679–685.

Steriade, M., Nuñez, A. & Amzica, F. (1993) Intracellular analysis of relations between the slow (< 1 Hz) neocortical oscillation and other sleep rhythms of the electroencephalogram. *Journal of Neuroscience*. 13 (8), 3266–3283.

Strange, K. (2008) Authorship: why not just toss a coin? *American Journal of Physiology - Cell Physiology*. 295 (3), C567–75.

Suner, S., Fellows, M.R., Vargas-Irwin, C., Nakata, G.K. & Donoghue, J.P. (2005) Reliability of signals from a chronically implanted, silicon-based electrode array in non-human primate primary motor cortex. *IEEE Trans Neural Syst Rehabil Eng*. 13 (4), 524–41.

Svoboda, K., Denk, W., Kleinfeld, D. & Tank, D.W. (1997) In vivo dendritic calcium dynamics in neocortical pyramidal neurons. *Nature*. 385 (6612), 161–165.

Swann, J.W., Al-Noori, S., Jiang, M. & Lee, C.L. (2000) Spine loss and other dendritic abnormalities in epilepsy. *Hippocampus*. 10 (5), 617–625.

Szabadics, J., Varga, C., Molnar, G., Olah, S., Barzo, P. & Tamas, G. (2006) Excitatory effect of GABAergic axo-axonic cells in cortical microcircuits. *Science*. 311 (5758), 233–235.

Takahashi, S., Anzai, Y. & Sakurai, Y. (2003) A new approach to spike sorting for multi-neuronal activities recorded with a tetrode—how ICA can be practical. *Neuroscience Research*. 46 (3), 265–272.

Thompson, S.A., Krishnan, B., Gonzalez-Martinez, J., Bulacio, J., Jehi, L., Mosher, J., Alexopoulos, A. & Burgess, R.C. (2015) Ictal infraslow activity in stereoelectroencephalography: Beyond the DC shift. *Clinical Neurophysiology*. doi: 10.1016/j.clinph.2015.03.020.

Thompson, S.M. & Gähwiler, B.H. (1989a) Activity-dependent disinhibition. I. Repetitive stimulation reduces IPSP driving force and conductance in the hippocampus in vitro. *J Neurophysiol*. 61 (3), 501–511.

Thompson, S.M. & Gähwiler, B.H. (1989b) Activity-dependent disinhibition. II. Effects of extracellular potassium, furosemide, and membrane potential on ECl- in hippocampal CA3 neurons. *J Neurophysiol*. 61 (3), 512–523.

Thompson, S.M. & Gähwiler, B.H. (1989c) Activity-dependent disinhibition. III. Desensitization and GABAB receptor-mediated presynaptic inhibition in the hippocampus in vitro. *J Neurophysiol*. 61 (3), 524–533.

Tian, G.-F., Azmi, H., Takano, T., Xu, Q., Peng, W., Lin, J., Oberheim, N., Lou, N., Wang, X., Zielke, H.R., Kang, J. & Nedergaard, M. (2005) An astrocytic basis of epilepsy. *Nat Med*. 11 (9), 973-81.

Timofeev, I. & Steriade, M. (2004) Neocortical seizures: initiation, development and cessation. *Neuroscience*. 123 (2), 299–336.

Timofeev, I., Grenier, F. & Steriade, M. (2004) Contribution of intrinsic neuronal factors in the generation of cortically driven electrographic seizures. *J Neurophysiol*. 92 (2), 1133–1143.

Timofeev, I., Grenier, F. & Steriade, M. (2002) The role of chloride-dependent inhibition and the activity of fast-spiking neurons during cortical spike-wave electrographic seizures. *Neuroscience*. 114 (4), 1115–1132.

Tolias, A.S., Ecker, A.S., Siapas, A.G., Hoenselaar, A., Keliris, G.A. & Logothetis, N.K. (2007) Recording Chronically From the Same Neurons in Awake, Behaving Primates. *J Neurophysiol*. 98 (6), 3780–3790.

Traub, R.D. & Miles, R. (1991) Multiple modes of neuronal population activity emerge after modifying specific synapses in a model of the CA3 region of the hippocampus. *Annals of the New York Academy of Sciences*. 627277–290.

Traub, R.D. & Wong, R.K. (1982) Cellular mechanism of neuronal synchronization in epilepsy. *Science*. 216 (4547), 745–747.

Trevelyan, A.J. & Schevon, C.A. (2013) How inhibition influences seizure propagation. *Neuropharmacology*. 69, 45–54.

Trevelyan, A.J., Bruns, W., Mann, E.O., Crepel, V. & Scanziani, M. (2013) The information content of physiological and epileptic brain activity. *J Physiol*. 591 (4), 799–805.

Trevelyan, A.J., Sussillo, D. & Yuste, R. (2007) Feedforward inhibition contributes to the control of epileptiform propagation speed. *Journal of Neuroscience*. 27 (13), 3383–3387.

Trevelyan, A.J., Sussillo, D., Watson, B.O. & Yuste, R. (2006) Modular propagation of epileptiform activity: evidence for an inhibitory veto in neocortex. *Journal of Neuroscience*. 26 (48), 12447–12455.

Truccolo, W., Donoghue, J.A., Hochberg, L.R., Eskandar, E.N., Madsen, J.R., Anderson, W.S., Brown, E.N., Halgren, E. & Cash, S.S. (2011) Single-neuron dynamics in human focal epilepsy. *Nature Neuroscience*. 14 (5), 635–641.

Tønnesen, J., Sørensen, A.T., Deisseroth, K., Lundberg, C. & Kokaia, M. (2009) Optogenetic control of epileptiform activity. *Proceedings of the National Academy of Sciences of the United States of America*. 106 (29), 12162–12167.

Vajda, I., van Pelt, J., Wolters, P., Chiappalone, M., Martinoia, S., van Someren, E. & van Ooyen, A. (2008) Low-frequency stimulation induces stable transitions in stereotypical activity in cortical networks. *Biophysical Journal*. 94 (12), 5028–5039.

Varela, F., Lachaux, J.P., Rodriguez, E. & Martinerie, J. (2001) The brainweb: phase synchronization and large-scale integration. *Nature Reviews Neuroscience*. 2 (4), 229–239.

Verzeano, M., Crandall, P.H. & Dymond, A. (1971) Neuronal activity of the amygdala in patients with psychomotor epilepsy. *Neuropsychologia*. 9 (3), 331–344.

Vigneswaran, G., Kraskov, A. & Lemon, R.N. (2011) Large Identified Pyramidal Cells in Macaque Motor and Premotor Cortex Exhibit 'Thin Spikes': Implications for Cell Type Classification. *Journal of Neuroscience*. 31 (40), 14235–14242.

Wadman, W.J. & Gutnick, M.J. (1993) Non-uniform propagation of epileptiform discharge in brain slices of rat neocortex. *Neuroscience*. 52 (2), 255–262.

Waziri, A., Schevon, C.A., Cappell, J., Emerson, R.G., McKhann, G.M. & Goodman, R.R. (2009) Initial surgical experience with a dense cortical microarray in epileptic patients undergoing craniotomy for subdural electrode implantation. *Neurosurgery*. 64 (3), 540–45.

Weinand, M.E., Carter, L.P., Patton, D.D., Oommen, K.J., Labiner, D.M. & Talwar, D. (1994) Long-term surface cortical cerebral blood flow monitoring in temporal lobe epilepsy. *Neurosurgery*. 35 (4), 657–664.

Wei, Y., Ullah, G. & Schiff, S.J. (2014) Unification of neuronal spikes, seizures, and spreading depression. *Journal of Neuroscience*. 34 (35), 11733–11743.

Weiss, S.A., Banks, G.P., McKhann, G.M., Goodman, R.R., Emerson, R.G., Trevelyan, A.J. & Schevon, C.A. (2013) Ictal high frequency oscillations distinguish two types of seizure territories in humans. *Brain*. 136 (12), 3796–3808.

Weiss, S.A., Lemesiou, A., Connors, R., Banks, G.P., McKhann, G.M., Goodman, R.R., Zhao, B., Filippi, C.G., Nowell, M., Rodionov, R., Diehl, B., McEvoy, A.W., Walker, M.C., Trevelyan, A.J., Bateman, L.M., Emerson, R.G. & Schevon, C.A. (2015) Seizure localization using ictal phase-locked high gamma: A retrospective surgical outcome study. *Neurology*. 84 (23), 2320–2328.

Wendling, F., Chauvel, P., Biraben, A. & Bartolomei, F. (2010) From Intracerebral EEG Signals to Brain Connectivity: Identification of Epileptogenic Networks in Partial Epilepsy. *Frontiers in Systems Neuroscience*. 4, 154.

Werhahn, K.J. (2010) Weakness and focal sensory deficits in the postictal state. *Epilepsy Behav*. 19 (2), 138–139.

Whittington, M.A., Cunningham, M.O., LeBeau, F.E.N., Racca, C. & Traub, R.D. (2010) Multiple origins of the cortical gamma rhythm. *Developmental Neurobiology*. 71 (1), 92–106.

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

Wiebe, S. & Jetté, N. (2012) Epilepsy surgery utilization: who, when, where, and why? *Current Opinion in Neurology*. 25 (2), 187–193.

Wiebe, S., Blume, W.T., Girvin, J.P., Eliasziw, M.Effectiveness and Efficiency of Surgery for Temporal Lobe Epilepsy Study Group (2001) A randomized, controlled trial of surgery for temporal-lobe epilepsy. *The New England Journal of Medicine*. 345 (5), 311–318.

Wilson, N.R., Runyan, C.A., Wang, F.L. & Sur, M. (2012) Division and subtraction by distinct cortical inhibitory networks in vivo. *Nature*. 488 (7411), 343–348.

Wong, B.Y. & Prince, D.A. (1990) The lateral spread of ictal discharges in neocortical brain slices. *Epilepsy Res*. 7 (1), 29–39.

Wong, M. (2005) Modulation of dendritic spines in epilepsy: Cellular mechanisms and functional implications. *Epilepsy & Behavior*. 7 (4), 569–577.

Wong, M. & Guo, D. (2013) Dendritic spine pathology in epilepsy: Cause or consequence? *Neuroscience*. 251 (C), 141–150.

Wood, F. & Black, M.J. (2008) A nonparametric Bayesian alternative to spike sorting. *Journal of Neuroscience Methods*. 173 (1), 1–12.

Wu, S., Kunhi Veedu, H.P., Lhatoo, S.D., Koubeissi, M.Z., Miller, J.P. & Lüders, H.O. (2014) Role of ictal baseline shifts and ictal high-frequency oscillations in stereo-electroencephalography analysis of mesial temporal lobe seizures. *Epilepsia*. 55 (5), 690–698.

Wykes, R.C., Heeroma, J.H., Mantoan, L., Zheng, K., MacDonald, D.C., Deisseroth, K., Hashemi, K.S., Walker, M.C., Schorge, S. & Kullmann, D.M. (2012) Optogenetic and potassium channel gene therapy in a rodent model of focal neocortical epilepsy. *Science Translational Medicine*. 4 (161), 161ra152.

Wykes, R.C., Kullmann, D.M., Pavlov, I. & Magloire, V. (2015) Optogenetic approaches to treat epilepsy. *Journal of Neuroscience Methods*. doi:10.1016/j.jneumeth.2015.06.004.

Wyler, A.R., Ojemann, G.A. & Ward, A.A. (1982) Neurons in human epileptic cortex: correlation between unit and EEG activity. *Annals of Neurology*. 11 (3), 301–308.

Xiong, Z.Q., Saggau, P. & Stringer, J.L. (2000) Activity-dependent intracellular acidification correlates with the duration of seizure activity. *Journal of Neuroscience*. 20 (4), 1290–1296.

Zhang, Z.J., Koifman, J., Shin, D.S., Ye, H., Florez, C.M., Zhang, L., Valiante, T.A. & Carlen, P.L. (2012) Transition to seizure: ictal discharge is preceded by exhausted presynaptic GABA release in the hippocampal CA3 region. *Journal of Neuroscience*. 32 (7), 2499–2512.

Ziburkus, J., Cressman, J.R., Barreto, E. & Schiff, S.J. (2006) Interneuron and pyramidal cell interplay during in vitro seizure-like events. *J Neurophysiol*. 95 (6), 3948–3954.