Newcastle University

School of Computing

Uncovering Patterns of Seizure Variability Within Individual Patients with Focal Epilepsy

Gabrielle Marie Schroeder

Submitted for the degree of Doctor of Philosophy in the School of Computing, Newcastle University

March 2022

Abstract

One of the primary challenges in treating epilepsy is that it is a dynamic disorder, with fluctuations in pathological brain activity and symptoms. Interestingly, seizures themselves can manifest in different ways, resulting in a range of spatiotemporal seizure evolutions within the same patient. Although such diversity may affect patient treatments and outcomes, the extent and characteristics of within-patient seizure variability are unknown. Here I investigated within-patient seizure variability by analysing seizure network evolutions in intracranial EEG (iEEG) recordings from patients with focal epilepsy. I first developed an approach for objectively comparing seizures and used my resulting "seizure dissimilarity" measure to characterise the extent and features of seizure variability. I then investigated how seizures changed over time in short-term recordings, revealing that variability in seizure evolutions could be explained by fluctuations over circadian and/or slower timescales. I next examined the relationship between seizure evolutions and seizure durations and found that these features could vary independently in the same patient: seizures with the same evolution could have dramatically different durations due to temporal "elasticity," and seizures with similar durations could have distinct evolutions. Finally, using chronic iEEG recordings, I explored how the occurrence and duration of different seizure networks changed over multiple timescales that were revealed by fluctuations in interictal spike rate. I found that all patients had seizure network feature(s) that were associated with specific spike rate fluctuations, suggesting that seizure evolutions were modulated over these timescales. My work provides the first extensive characterisation of variability in within-patient seizure evolutions as well as a framework for quantitatively comparing seizures. Moreover, my results suggest that various modulatory factors, operating over different timescales, influence seizure evolutions. Understanding the mechanisms that shape seizure features within the same brain could provide new opportunities

for controlling the full repertoire of seizures in each patient.

Declaration

I declare that this thesis is my own work unless otherwise stated. No part of this thesis has previously been submitted for a degree or any other qualification at Newcastle University or any other institution.

Gabrielle Marie Schroeder

2 March 2022

Publications

Portions of the work within this thesis have been documented in the following publications:

Most of the content in Chapters 2, 3, and 4 has been previously published in

Schroeder, G. M., Diehl, B., Chowdhury, F.A., Duncan, J. S., de Tisi, J., Trevelyan, A. J., Forsyth, R., Jackson, A., Taylor, P. N., and Wang, Y. (2020). Seizure pathways change on circadian and slower timescales in individual patients with focal epilepsy. *Proceedings of the National Academy of Sciences*, 117(20):11048-11058.

I performed all of the data analysis, created all of the data visualisations, and wrote all of the content from this publication that is included in my thesis, with input from my supervisors and other co-authors. Data collection for the data that was not publicly available was performed by our collaborators and co-authors at University College London Hospital.

Chapters 5 and 6 are available as preprints and have been submitted for publication:

Schroeder, G. M., Chowdhury, F.A., Cook, M. J., Diehl, B., Duncan, J. S., Karoly, P. J., Taylor, P. N., Wang, Y. (2021). Seizure pathways and seizure durations can vary independently within individual patients with focal epilepsy. *arXiv:q-bio.NC*, 2109.06672

Schroeder, G. M., Karoly, P. J., Maturana, M., Taylor, P. N., Cook, M. J., Wang, Y. (2022). Chronic iEEG recordings and interictal spike rate reveal multiscale temporal modulations in seizure states. *arXiv:q-bio.NC*, 2201.11600

Unless otherwise specified, I performed all of the data analysis, created all of the data visualisations, and wrote all of the content for these chapters with input from my supervisors

and other coauthors. Our collaborators and coauthors at Melbourne University obtained, extracted, and partially preprocessed the NeuroVista seizure data and also computed the interictal spike rate data analysed in Chapter 6.

Outside of my thesis work, I have contributed to the following published papers and preprints:

Taylor, P. N., Papasavvas, C. A., Owen, T. W., **Schroeder, G. M.**, Hutchings, F. E., Chowdhury, F. A., Diehl, B., Duncan, J. S., McEvoy, A. W., Miserocchi, A., de Tisi, J., Vos, S. B., Walker, M. C., Wang, Y. (2022). Normative brain mapping of interictal intracranial EEG to localize epileptogenic tissue. *Brain*, DOI: 10.1093/brain/awab380

Panagiotopoulou, M., Papasavvas, C. A., **Schroeder, G. M.**, Thomas, R. H., Taylor, P. N., Wang, Y. (2022). Fluctuations in EEG band power at subject-specific timescales over minutes to days explain changes in seizure evolutions. *Human Brain Mapping*, DOI: 10.1002/hbm.25796

Sinha, N., Peternell, N., **Schroeder, G. M.**, de Tisi, J., Vos, S. B., Winston, G. P., Duncan, J. S., Wang, Y., Taylor, P. N. (2021). Focal to bilateral tonic-clonic seizures are associated with widespread network abnormality in temporal lob epilepsy. *Epilepsia*, 62(3):729-741

Papasavvas, C. A., **Schroeder, G. M.**, Diehl, B., Baier, G., Taylor, P. N., Wang, Y. (2020). Band power modulation through intracranial EEG stimulation and its cross-session consistency. *Journal of Neural Engineering*, 17(5):054001

Wang, Y., Sinha, N., **Schroeder, G. M.**, Ramaraju, S., McEvoy, A. W., Miserocchi, A., de Tisi, J., Chowdhury, F. A., Diehl, B., Duncan, J. S., Taylor, P. N. (2020). Interictal intracranial electroencephalography for predicting surgical success: The importance of space and time. *Epilepsia*, 61(7):1417-1426

Wang, Y., **Schroeder, G. M.**, Sinha, N., Taylor, P. N. (2019). Personalised network modelling in epilepsy. *arXiv:q-bio.NC*, 1901.01024v1

Acknowledgements

First, I would like to thank my PhD supervisors, Yujiang Wang and Andrew Jackson, for their guidance, support, and mentorship. I am especially grateful to Yujiang Wang for accepting me as one of her first PhD students and helping me return to Newcastle for my PhD. Throughout my PhD, I have also benefited from interactions with and support from the Computational Neurology, Neuroscience, and Psychiatry Lab, the School of Computing, and our clinical and experimental neuroscience collaborators at Newcastle University. In particular, I would like to thank Peter Taylor, Rob Forsyth, and Andrew Trevelyan for their input on this work. I am also grateful to the School of Computing for funding my PhD.

The data used in my projects comes from both open source datasets and our collaborators. I am grateful to the IEEG Portal and our colleagues at University College London Hospital (UCLH) for the epilepsy monitoring unit data and long-term canine recordings. My work has also benefited from discussions with my UCLH co-authors, Beate Diehl, John Duncan, Jane de Tisi, and Fahmida Chowdhury. The last chapter of my thesis would not have been possible without a unique dataset, and I am grateful to our collaborators at Melbourne University, Mark Cook, Philippa Karoly, and Matias Maturana, for sharing the NeuroVista seizure and spike rate data and for our discussions on this project.

Finally, I would like to thank my partner, James Skelton, for his support throughout my PhD, especially during the last two strange and unpredictable years of the COVID pandemic.

Contents

List of Figures

List of Tables

Acronyms

AUC area under the curve

ARI adjusted Rand index

AED antiepileptic drug

CEEMDAN complete ensemble empirical mode decomposition with adaptive noise

DTW dynamic time warping

EEG electroencephalographic

EMD empirical mode decomposition

EMU epilepsy monitoring unit

FDR false discovery rate

iEEG intracranial EEG

ILAE International League Against Epilepsy

MDS multidimensional scaling

MSE mean squared error

NMF non-negative matrix factorisation

PLV phase locking value

UCLH University College London Hospital

UPGMA unweighted pair group method with arithmetic mean

Chapter 1. Background: Brain Dynamics in Focal Epilepsy

Contents

1.1 Focal Epilepsy

Epilepsy is a neurological disorder in which pathological neural dynamics lead to spontaneous, recurrent seizures (Badawy et al., 2012). In focal epilepsies, a patient-specific epileptogenic zone triggers the onset of seizure activity, which can subsequently recruit neighbouring healthy regions to the seizure (Rosenow and Lüders, 2001). Although the epileptogenic zone is localised, altered interactions between distant brain regions are also thought to create an "epileptic network" that further contributes to the pathology (Bernhardt et al., 2015; Kramer and Cash, 2012; Spencer, 2002). Additionally, widespread alterations in the balance between neuronal excitation and inhibition (Badawy et al., 2012) play an important role in producing the pathological brain activity observed in epilepsy. Uncovering mechanisms that contribute to the development of epilepsy (epileptogenesis) and seizure occurrence (ictogenesis) is an active area of research, with researchers and clinicians investigating a range of potential factors such as genetic (Myers and Mefford, 2015), circuit (Paz and Huguenard, 2015), and structural network (Bernhardt et al., 2015; Englot et al., 2016; Sinha et al., 2021a,b) abnormalities.

Focal epilepsy is commonly treated with antiepileptic drugs (AEDs) that cause widespread changes to neural exciability (Meisel et al., 2016, 2015) through molecular targets such as ion channels and neurotransmitter receptors (Macdonald and Kelly, 1995). In those patients that continue experiencing seizures despite medication, surgical treatment to resect the epileptogenic tissue is often pursued (Rosenow and Lüders, 2001); however, in 50-70% of patients, surgery fails to completely abolish seizures or only provides short-term relief (de Tisi et al., 2011). Due to the limitations of these approaches, researchers and clinicians are developing novel treatment strategies such as brain stimulation (Jarosiewicz and Morrell, 2021; Sisterson et al., 2019; Wang et al., 2015) and chronotherapy, in which AED doses depend on fluctuations in seizure risk (Baud and Rao, 2018; Ramgopal et al., 2013). To improve existing treatment strategies and develop new treatment options, a better understanding of the mechanisms underlying epilepsy and seizures is needed.

One of the challenges in treating epilepsy is that seizure features such as patterns and spread vary, even in patients with the same type of epilepsy (Alarcon et al., 1995; Fisher

et al., 2017; Jiménez-Jiménez et al., 2015; Martinet et al., 2015; Saggio et al., 2020; Salami et al., 2021, 2020). This diversity in how the disease manifests suggests that the underlying mechanisms of seizure processes also vary. As such, patients may require tailored treatment strategies (Baud and Rao, 2018; Freestone et al., 2017; Sisterson et al., 2019; Wang et al., 2015, 2019). Indeed, surgical outcome is associated with seizure onset patterns (Alarcon et al., 1995; Jiménez-Jiménez et al., 2015; Lagarde et al., 2019) and propagation patterns (Martinet et al., 2015), though the reasons certain characteristics improve treatment outcomes is poorly understood. Thus, relating seizure features to their underlying mechanisms and using that knowledge to guide clinical treatments will be key for achieving seizure freedom in such a heterogeneous group of patients.

1.2 Capturing and Analysing Epileptic Brain Activity

Many researchers use computational approaches to explore mechanisms underlying epilepsy and seizures (Baud et al., 2018; Jirsa et al., 2014; Karoly et al., 2018a; Leguia et al., 2021; Proix et al., 2018; Wang et al., 2017) and develop new treatment strategies (Goodfellow et al., 2016; Karoly et al., 2017; Kini et al., 2019; Proix et al., 2021; Sinha et al., 2016; Stirling et al., 2021; Wang et al., 2015, 2019). Much of this research involves analysing recordings of neural activity from patients with epilepsy, such as electroencephalographic (EEG) recordings, and this is the approach I will take in my thesis. In particular, the data in this work comes from intracranial EEG (iEEG) recordings, in which implanted subdural and/or depth electrodes record activity from either the cortical surface or deeper neural tissue, respectively (Rosenow and Lüders, 2001). These recordings are often performed as part of pre-surgical monitoring in epilepsy monitoring units (EMUs) to collect data for localising epileptogenic tissue prior to resective surgery (Rosenow and Lüders, 2001). For each patient, epilepsy monitoring unit (EMU) iEEG recordings typically last a few days to a few weeks and thus provide a snapshot of the patient's longer-term brain dynamics (Baud et al., 2021). More recently, seizure prediction and neurostimulation studies have obtained chronic (multiple months to years) iEEG recordings in human patients with focal epilepsy (Cook et al., 2013; Jarosiewicz and Morrell, 2021) and animals with focal-onset seizures (Howbert et al., 2014). Chronic iEEG recordings can capture slower timescales of

brain dynamics (Baud et al., 2018; Baud and Rao, 2018; Karoly et al., 2018a, 2021; Leguia et al., 2021) and rarer events (King-Stephens et al., 2015) compared to EMU recordings. Both EMU and chronic iEEG recordings capture seemingly normal background activity as well as pathological events such as interictal spikes, high frequency oscillations, and seizures (Baud et al., 2018; Chen et al., 2021; Karoly et al., 2016, 2018a; Leguia et al., 2021; Ung et al., 2016).

Fundamentally, iEEG signals are a type of multivariate time series that describe spatial changes in brain activity at a high temporal resolution, and thus can be analysed using a variety of time series analysis methods. In the following sections, I will briefly discuss computational approaches that are commonly applied to this data.

1.2.1 Univariate times series analysis

Univariate times series measures describe features of a single time series and can capture properties such as signal amplitude, variance, and spectral composition (Fulcher, 2017). When analysing iEEG recordings, univariate measures can either be applied to the signal from a single electrode (also known as a recording channel), or individually to all electrodes. In the latter case, computing a univariate measure of an iEEG signal with n recording electrodes yields n features of the iEEG that describe the spatial pattern of the measure across the recording layout. Multiple univariate measures can also be computed for the same data to characterise different properties of the neural signal.

Often, researchers are interested in how brain activity changes over time during interictal (between seizure), preictal (before seizure), ictal (seizure), and postictal (after seizure) periods. Therefore, to describe the temporal evolution of brain activity, the iEEG recording can be divided into smaller segments. A common approach is to apply a "sliding window," in which the signal is divided into successive, equivalent-sized segments (Karoly et al., 2018b; Panagiotopoulou et al., 2022). Neighbouring windows can overlap in time to further increase the measure's temporal resolution. The univariate measure is then applied to the time series of each segment and recording electrode. If the recording is divided into T windows, this analysis produces a new $n \times T$ time series that describes the spatiotemporal evolution of the measure during the selected time period.

A number of univariate measures have been applied to iEEG recordings to understand interictal brain dynamics (Maturana et al., 2020; Panagiotopoulou et al., 2022; Taylor et al., 2021), characterise seizure patterns and evolutions (Jiménez-Jiménez et al., 2015; Salami et al., 2020; Wendling et al., 1997, 1996, 1999; Wu and Gotman, 1998), and predict seizure occurrence and features (Howbert et al., 2014; Maturana et al., 2020; Naftulin et al., 2018). One common measure is signal band power in specific frequency bands (Howbert et al., 2014; Panagiotopoulou et al., 2022; Taylor et al., 2021), which are defined based on past observations of healthy and pathological brain rhythms (Blum and Rutkove, 2007). While the exact definitions of these frequency bands can slightly vary, they are generally approximately divided into delta (1-4 Hz), theta (4-8 Hz), alpha (8-13 Hz), beta (13-30 Hz), gamma (30-80 Hz), and high gamma (80-150 Hz) (Groppe et al., 2013). This spectral analysis captures normal spatial variability in the frequency of brain activity (Blum and Rutkove, 2007; Groppe et al., 2013; Taylor et al., 2021) as well as pathological ictal (Alarcon et al., 1995; Blum and Rutkove, 2007; Jiménez-Jiménez et al., 2015; Salami et al., 2020; Wendling et al., 1997, 1996, 1999; Wu and Gotman, 1998) and interictal (Blum and Rutkove, 2007; Chen et al., 2021; Gliske et al., 2018) rhythms.

Univariate time series analysis can also be used to automatically detect and characterise brief pathological interictal activity such as epileptiform spikes (Karoly et al., 2016) and high frequency oscillations (Gliske et al., 2018) in each recording channel. While these events can be detected visually on iEEG traces, marking them in long recordings is prohibitively time-consuming. By automatically detecting abnormal interictal activity, researchers have been able to characterise how their spatial patterns (Chen et al., 2021; Gliske et al., 2018) and/or frequency (Baud et al., 2018; Karoly et al., 2016; Leguia et al., 2021) change over time and relate to other features such as seizures.

1.2.2 Functional networks

While a univariate analysis captures properties of individual brain areas, interactions between different parts of the brain play a crucial role in epileptic processes (Kramer and Cash, 2012; Spencer, 2002). To describe these "network" interactions, many researchers compute functional networks, also known as functional connectivity, from iEEG record-

ings. In this scenario, each node represents a brain region whose activity has been captured by an electrode. The network edges correspond to the statistical relationship (e.g., correlation or coherence) between the iEEG signals of each pair of recorded brain regions. Thus, an iEEG recording with n electrodes can be transformed into a $n \times n$ functional network, with each entry corresponding to the pairwise interaction of two recordings sites. As with univariate measures, time-varying functional networks can be computed from windowed iEEG recordings to explore changes in brain dynamics (Burns et al., 2014; Khambhati et al., 2017, 2015; Kramer et al., 2010; Mitsis et al., 2020; Schindler et al., 2007b). The spatiotemporal evolution of functional connectivity can be quantified by identifying recurring patterns of interactions (Burns et al., 2014; Frusque et al., 2020; Khambhati et al., 2017) and tracking changes in node-wise or global network measures (Burns et al., 2014; Khambhati et al., 2016; Kramer et al., 2010; Meisel et al., 2015; Mitsis et al., 2020; Schindler et al., 2007b). Such work has contributed to our knowledge of how network interactions change during seizures and relate to the underlying pathology. For example, functional network structure has been linked to properties such as seizure spread (Khambhati et al., 2016). Further, network interactions during seizures may reveal the epileptogenic zone (Burns et al., 2014; Goodfellow et al., 2016; Kini et al., 2019), demonstrating the potential usefulness of network analysis for clinical applications.

1.2.3 Computational modelling

The above analytical techniques can be complemented by computational modelling of neural activity (Jirsa et al., 2014; Proix et al., 2018; Wang et al., 2019, 2017; Wendling et al., 2002). A biophysical model of the brain contains interacting components that correspond to different neurons (Hodgkin and Huxley, 1952; Izhikevich, 2003) or populations of neurons (Wang et al., 2019; Wilson and Cowan, 1972). The model then simulates neural activity from different components, with differential equations describing how the activity changes over time. Model parameters, which describe features and interactions of the model components, can be changed to explore how such alterations affect brain activity (Jirsa et al., 2014; Wang et al., 2017). Alternatively, model parameters can be derived from neural data to elucidate the possible underlying interactions between neural

populations (Karoly et al., 2018b). Researchers have used biophysical models to explore mechanisms involved in processes such as seizure onset, propagation, and termination (Jirsa et al., 2014; Kramer et al., 2012; Proix et al., 2018; Wang et al., 2017; Wendling et al., 2002). Through the incorporation of patient-specific data, these models also provide opportunities to test treatment strategies (Goodfellow et al., 2016; Sinha et al., 2016; Wang et al., 2019) and propose novel interventions, such as neurostimulation (Wang et al., 2015).

1.3 Seizure Variability Within Individual Patients with Focal Epilepsy

As discussed earlier, many seizure features can vary between patients (Fisher et al., 2017; Proix et al., 2018; Saggio et al., 2020; Salami et al., 2021), and this diversity may contribute to differences in treatment responses (Alarcon et al., 1995; Jiménez-Jiménez et al., 2015; Lagarde et al., 2019; Martinet et al., 2015). On the other hand, seizures in the same patient are often described as stereotyped due to conserved EEG evolutions (Burns et al., 2014; Karoly et al., 2018b; Wagner et al., 2015) and symptoms (Chauvel and McGonigal, 2014; McGonigal, 2020). Nonetheless, certain features such as seizure onset patterns (Jiménez-Jiménez et al., 2015; Salami et al., 2020) and durations (Cook et al., 2016) can vary from seizure to seizure within the same patient. What features of intra-patient seizures vary, the mechanisms that support this variation, and its implications for treatment remain open questions. In part, these questions have been unanswered due to the challenges of objectively and quantitatively comparing the complex spatiotemporal features of seizures. In the following sections, I review current approaches for comparing characteristics of within-patient seizures and past findings of within-patient variability using these frameworks.

1.3.1 ILAE clinical seizure types

ILAE clinical seizure types provide a broad classification of seizures based on their extent of spread and certain clinical symptoms (Fisher et al., 2017). These clinical types do not group seizures based on underlying mechanisms or seizure EEG patterns; rather, they are

defined as "a useful grouping of seizure characteristics for purposes of communication in clinical care, teaching, and research" (Fisher et al., 2017).

In this classification, focal-onset seizures are divided into two main categories (Fisher et al., 2017):

1. **Focal:** Focal seizures are defined as originating in the network of one hemisphere and remain in that hemisphere during their ensuing propagation.

2. **Focal to bilateral tonic-clonic (formerly "secondarily generalised"):** Focal to bilateral tonic-clonic seizures begin in one hemisphere, but spread to the contralateral hemisphere, thus engaging bilateral networks and leading to tonic and clonic contractions (Noachtar and Peters, 2009). For consistency with my clinical metadata, I will use the older term "secondarily generalised" to refer to these seizures throughout this thesis.

A seizure can also be classified as "subclinical" or "electrographic only" if it appears on the EEG recording, but does not have any corresponding clinical symptoms (Farooque and Duckrow, 2014). While focal seizures can be further subdivided based on their salient or early clinical symptoms (Fisher et al., 2017), computational research of ILAE clinical seizure types has largely focused on subclinical, focal, and secondarily generalised seizures (Badawy et al., 2009; Karthick et al., 2018; Khambhati et al., 2016; Naftulin et al., 2018; Sinha et al., 2021a), and I also focus on those distinctions in this thesis.

It is well known that the same patient often experiences multiple clinical seizure types (Farooque and Duckrow, 2014; Karthick et al., 2018; Marciani and Gotman, 1986; Naftulin et al., 2018). Due to the severity of secondarily generalised seizures, there has been particular interest in understanding why some seizures remain focal while others secondarily generalise. As such, research on secondary generalisation provides some of the most concrete examples of how spatiotemporal brain dynamics may influence seizure features within and between patients. In particular, widespread functional network (Khambhati et al., 2016), band power (Karthick et al., 2018; Naftulin et al., 2018), and cortical excitability (Badawy et al., 2009; Enatsu et al., 2012) patterns during the preictal period

and/or seizure onset have been implicated in whether seizures secondarily generalise. These findings suggest that brain state during and before seizure onset plays a crucial role in shaping the extent of seizure spread.

Importantly, the ILAE clinical seizure classification does not distinguish seizures by their specific propagation patterns, EEG rhythms, or other characteristics (Fisher et al., 2017). Thus, this seizure classification provides limited insight into different spatiotemporal seizure features.

1.3.2 Seizure semiology

Beyond the symptoms highlighted by the broad ILAE clinical seizure classification, seizures can cause a number of different symptoms, and a seizure's set of clinical symptoms are referred to as its semiology (Noachtar and Peters, 2009). Seizure symptoms fall into one of four categories (sensorial, motor, consciousness, or autonomic) depending on what system they affect; for example, auras are subjective symptoms, such as auditory hallucinations or a fearful feeling, that belong to the sensorial sphere (Noachtar and Peters, 2009; Rossetti and Kaplan, 2010). A seizure's semiology often consists of multiple symptoms from multiple categories (Noachtar and Peters, 2009), and the sequence of symptoms is often consistent across a patient's seizures (Chauvel and McGonigal, 2014; McGonigal, 2020). This stereotyped semiology arises because seizure symptoms are linked to which areas are activated by the seizure (Noachtar and Peters, 2009); thus, seizures with similar evolutions will activate the same areas and produce similar symptoms (Chauvel and McGonigal, 2014). Inferring seizure spread and localising seizure onset based on semiology is an active area of research (Chauvel and McGonigal, 2014; McGonigal, 2020; Rossetti and Kaplan, 2010). However, concrete mappings between clinical symptoms, brain regions, and brain network interactions have not been established (Chauvel and McGonigal, 2014; McGonigal, 2020), and the localisation reliability and specificity of different symptoms varies (McGonigal, 2020; Rossetti and Kaplan, 2010).

Using semiology alone to classify seizures has several limitations. First, seizure spread to a brain region does not necessarily cause symptoms, and thus a seizure's full propagation cannot be inferred from semiology alone (Noachtar and Peters, 2009). Further, within

the same patient, seizures with distinct onsets can have the same semiology, likely due to shared networks engaged by both types of seizures (Vaugier et al., 2017). As such, semiology does not accurately distinguish seizures with different evolutions. Like ILAE clinical types, although certain symptoms are associated with specific EEG patterns (Noachtar and Peters, 2009) and the frequency of seizure activity can shape symptoms (Chauvel and McGonigal, 2014; McGonigal, 2020), semiology does not explicitly distinguish between seizure spread with different features. Finally, a seizure's semiology can be difficult to determine and usually requires an experienced observer to classify the symptoms (McGonigal, 2020). During presurgical monitoring, clinicians use video-EEG to review a patient's seizure semiology (Rosenow and Lüders, 2001), but such data is not available for chronic iEEG recordings. Even with footage of the seizure, some symptoms such as loss of consciousness can be difficult to evaluate (Ali et al., 2012; McGonigal, 2020) and semiology assessments can vary between clinically trained observers (Benbir et al., 2013).

1.3.3 Seizure EEG patterns

Seizures are also characterised by different patterns of activity on EEG that can be classified by their frequency, amplitude, and temporal evolutions (Blum and Rutkove, 2007). Many studies have classified inter- and intra-patient seizures by their onset patterns, such as low-voltage fast activity or high amplitude spikes (Alarcon et al., 1995; Jiménez-Jiménez et al., 2015; Lagarde et al., 2019; Salami et al., 2020). Different onset patterns appear to have different underlying mechanisms (Wang et al., 2017) and are associated with different surgical outcomes (Alarcon et al., 1995; Jiménez-Jiménez et al., 2015; Lagarde et al., 2019); thus, onset patterns can provide important information about a patient's pathology. While onset patterns are often consistent from seizure to seizure in a given patient, some patients have multiple onset patterns (Jiménez-Jiménez et al., 2015; Salami et al., 2020), suggesting that different mechanisms for seizure onset may co-exist in the same brain.

Variability in seizure EEG patterns beyond seizure onset have been less well-explored, although some studies have examined variability in seizure termination features, such as bursting patterns (Salami et al., 2021) and whether termination is synchronous across

electrodes (Afra et al., 2015; Salami et al., 2021). As with seizure onset, the same patient can have multiple termination patterns (Afra et al., 2015; Salami et al., 2021), indicating that this feature is also not stereotyped within patients. Seizures of the same ILAE clinical seizure type can also have different termination patterns (Salami et al., 2021), further demonstrating that clinical types do not capture many spatiotemporal aspects of seizure variability.

While these classifications provide useful insights into mechanisms of seizure onset and termination, they are not designed to quantitatively compare seizures in the same patient. In particular, focusing solely on seizure onset and/or termination misses possible differences in other parts of seizure evolutions. Additionally, these classifications do not account for spatial differences between seizures, such as different onset locations or which channels terminate early in an asynchronous termination.

1.3.4 Dynamical classes of seizure onset and offset

Seizures can be modelled using a dynamical system that qualitatively changes over different timescales due to underlying factors (Jirsa et al., 2014). Seizures can then be characterised by the qualitative transitions, or bifurcations, that the model undergoes at seizure onset and termination, which can be identified from their different iEEG signatures (Jirsa et al., 2014). Recently, Saggio et al. (2020) used this framework to describe seizure "dynamotypes" in patients with focal epilepsy, with each dynamotype corresponding to a set of seizure onset and offset bifurcations. They found that individual patients can have multiple dynamotypes, and this co-existence of dynamotypes can be explained by the proximity of different bifurcations in a computational model of neural activity. While they focused on classifying seizure onset and offset dynamics, Saggio et al. (2020) also demonstrated that seizure evolutions can be described as pathways through this computational model, explaining more complex patterns of seizure onset, evolution, and termination. This classification provides a conceptual and mathematical framework for understanding inter- and intra-patient seizure diversity and suggests universal dynamical mechanisms for seizure onsets and terminations. However, the approach also has limitations: some empirical seizure recordings can be difficult to classify, not all recordings

contain sufficient information to distinguish certain dynamotypes, and the classification does not take into account spatial differences in seizure evolutions.

1.3.5 Spatiotemporal seizure evolutions

As discussed above, most seizure comparisons based on EEG focus on specific parts of seizures and do not capture spatial differences in seizure activity. However, several studies have proposed measures for comparing spatiotemporal seizure evolutions, defined as the spatial changes in EEG features over the course of a seizure. These approaches therefore first describe seizure evolutions as multivariate time series that capture the time-varying spatial patterns of specific seizure features (Burns et al., 2014; Dorr et al., 2007; Le Bouquin-Jeannès et al., 2002; Wendling et al., 1996, 1999; Wu and Gotman, 1998). The evolutions of a pair of seizures can then be compared using time series analysis methods, ultimately yielding a single number for the (dis)similarity of the two seizures. Thus, unlike other methods, this approach quantifies the pairwise similarities between a patient's seizures.

Many of these studies described seizure evolutions by computing multiple time-varying univariate measures (Wendling et al., 1996, 1999; Wu and Gotman, 1998) or time-varying functional connectivity (Dorr et al., 2007; Le Bouquin-Jeannès et al., 2002) and then computing a "distance" between each pair of seizure evolutions. Specifically, each channel's evolution was described by a sequence of states that summarised the channel's changing properties, and a seizure's evolution was then described by the time-varying states of all channels. The studies then compared pairs of within-patient seizures by computing the edit distance, or the cost of transforming one sequence into another by inserting, deleting, or switching states, of the multivariate state time series. Interestingly, Wu and Gotman (1998) observed variability in within-patient seizure evolutions, with some patients having a range of seizure edit distances and groupings of similar seizures. Importantly, the edit distance measure must be modified (e.g., as by Dorr et al. (2007)) to disentangle differences in seizure evolutions and seizure durations; otherwise, the insertion cost penalises seizures that take different amounts of time to complete the same evolution. Approaches that recognise similar seizure evolutions, even if the seizure progress at different rates, are

needed to characterise variability in different seizure features. Additionally, these past studies were limited to a small number of patients and/or seizures per patient, limiting their insights into the prevalence and characteristics of within-patient seizure variability.

More recently, Burns et al. (2014) took a different approach to comparing seizure evolutions. In this study, the authors instead clustered each patient's seizure functional networks into a small number of states that summarised different network patterns. Each seizure could then be described as a progression of these network states. Each seizure's dynamics were further characterised by a transition matrix containing transition probabilities between the network states. Burns et al. (2014) found that these transition matrices were consistent across seizures in half of their cohort of 12 patients, indicating similar within-patient seizure network evolutions. The variable transition matrices in the other patients suggests that seizure network evolutions can also differ across patients; however, the variability was not explored in this study. While Burns et al. (2014) explored the overall consistency of within-patient transition matrices, the same general approach could be used to create a pairwise measure of seizure similarity. The main disadvantages of this framework would be that 1) state descriptions of seizure evolutions may miss more subtle differences between seizures, and 2) seizure transition matrices do not preserve the order of state occurrence if any states occur multiple times in the same seizure. As such, state transition matrices do not fully describe seizure evolutions.

1.4 Thesis Aims and Contributions

Epilepsy is a dynamic disorder, and recent studies have highlighted the importance of understanding how pathological brain dynamics fluctuate over time within individual patients (Baud et al., 2018; Chen et al., 2021; Gliske et al., 2018; Karoly et al., 2021; Leguia et al., 2021; Mitsis et al., 2020; Saggio et al., 2020). While many studies have explored variability in interictal events and/or seizure occurrence (Baud et al., 2018; Chen et al., 2021; Gliske et al., 2018; Karoly et al., 2021; Leguia et al., 2021), little is known about variability in seizures themselves. In this thesis, I addressed this knowledge gap by answering three primary open questions:

1. How prevalent is variability in seizure evolutions within individual patients?

2. Does a seizure's evolution also determine its duration, or is there independent variability in these features?

3. How do seizure features change over different timescales?

My research chapters make the following contributions towards these aims:

- In Chapter 3, I first developed a new approach for objectively comparing within-patient seizure evolutions. I then used my pairwise "seizure dissimilarity" measure to compare seizure network evolutions in EMU iEEG recordings of patients with focal epilepsy and demonstrated that all patients had seizure variability. I additionally used my measure to (1) explore whether within-patient seizures could be clustered into distinct groups with different evolutions and (2) determine whether seizure variability was explained by differences in ILAE clinical seizure types.

- In Chapter 4, I investigated how seizures change over time in the same cohort of EMU patients. Interestingly, more similar seizures tended to occur closer together in time within each patient, and within-patient variability in seizure evolutions could be explained by fluctuations over circadian and/or slower timescales. These results suggest that underlying, time-varying factors influence seizure evolutions, producing the observed within-patient seizure variability.

- In Chapter 5, I next examined the relationship between seizure evolutions and seizure durations in EMU iEEG recordings as well as chronic iEEG recordings in human patients with focal epilepsy and canines with focal-onset seizures. I found that these features could vary independently in the same patient, indicating that seizures can have separate axes of variability in different features. Most notably, seizures with the same evolution could have dramatically different durations due to temporal "elasticity" in seizure evolutions, and seizures with similar durations could have distinct evolutions.

- Finally, in Chapter 6, I used the insights from the previous chapters to explore how within-patient seizure features varied over longer timescales in chronic iEEG recordings of human patients with focal epilepsy. I found that the occurrence and duration of different seizure networks could change over multiple timescales that were revealed by fluctuations in interictal spike rate. These results provide stronger evidence that seizure features are modulated over multiple timescales within individual patients.

In Chapter 7, I discuss extensions for this work, including mapping within-patient seizure variability, uncovering mechanisms underlying changes in seizure features, and determining the clinical implications of fluctuating seizure evolutions.

Chapter 1: Background: Brain Dynamics in Focal Epilepsy

Chapter 2. Subject Data and Computation of Seizure Network Evolutions

Contents

In this chapter, I describe the data and computational methods that are used throughout the subsequent research chapters of my thesis. In particular, I describe how I compute seizure network evolutions from seizure intracranial EEG (iEEG) recordings, reduce the dimensionality of these evolutions and compute progressions of seizure network states using non-negative matrix factorisation (NMF), visualise seizure network evolutions using multidimensional scaling (MDS), and compare within-subject network evolutions using dynamic time warping (DTW).

2.1 Subjects and Subject Seizure Data

In this work, we analysed iEEG data from three cohorts of subjects:

- **epilepsy monitoring unit (EMU) patients:** Thirty-one patients with drug-resistant focal epilepsy recorded in EMUs as part of presurgical identification of their epileptogenic zones (data used Chapters 3, 4, and 5).
- **NeuroVista patients**: Ten patients with drug-resistant focal epilepsy who were chronically implanted with iEEG electrodes as part of a NeuroVista seizure prediction study (data used in Chapters 5 and 6).
- **Canine subjects**: Three dogs with focal-onset seizures who were chronically implanted with iEEG electrodes as part of a NeuroVista seizure prediction study (data used in Chapter 5).

All of research chapters in this thesis were retrospective studies that did not influence data collection or patient treatments. Details on each cohort are below.

2.1.1 Epilepsy monitoring unit patients

We analysed seizures from 13 patients from the Mayo Clinic and the Hospital of the University of Pennsylvania (available on the IEEG Portal, www.ieeg.org (Kini et al., 2016; Wagenaar et al., 2013) and 18 patients from the University College London Hospital (UCLH) who were diagnosed with refractory focal epilepsy and underwent presurgical monitoring. Patients were selected without reference to the cause or other characteristics of their

pathology. All IEEG Portal patients gave consent to have their anonymised iEEG data publicly available on the International Epilepsy Electrophysiology Portal (www.ieeg.org) (Kini et al., 2016; Wagenaar et al., 2013). For the UCLH patients, their iEEG was anonymised and exported, and the anonymised data was subsequently analysed in this study under the approval of the Newcastle University Ethics Committee (reference number 6887/2018).

For each patient, the placement of the intracranial electrodes was determined by the clinical team, independent of this study. Ictal segments that were at least 10s long were identified and extracted for the analysis based on clinical seizure markings. To be included in the study, each patient was required to have had at least six seizures suitable for the analysis. This threshold was chosen to allow examination of seizure variability in a broad cohort of patients, while still ensuring that enough seizures were observed to draw conclusions about the characteristics of seizure variability in each patient. Seizures were excluded from the analysis if they did not have clear electrographic correlates (with clear onset and termination), if they were triggered by/occurred during cortical stimulation, if they had noisy segments, or if they had large missing segments. Periods of status epilepticus and continuous epileptiform discharges were also excluded. However, electrographic seizures without clinical correlates (i.e., subclinical seizures) were included in the analysis, as they may have either similar or disparate dynamics (relative to clinical seizures) that convey clinically relevant information (Farooque and Duckrow, 2014). Additional information about each patient and the analysed seizures is shown in Table 2.1.

2.1.2 Canine subjects

To explore seizure variability on longer timescales and in non-human subjects, iEEG was also analysed from three canine subjects with focal-onset seizures due to naturally occurring epilepsy that underwent prolonged recordings as part of a seizure prediction study (Howbert et al., 2014) (recording data available on the IEEG Portal, www.ieeg.org (Kini et al., 2016; Wagenaar et al., 2013)). Metadata for the canine subjects is provided in Table 2.2. As for the EMU patients, canine seizures were included in the analysis if were at least 10s long and lacked noisy or missing segments. Seizure onset was determined

Subject	Hospital	Age (yrs)	Sex	Hemisphere	Lobe	Pathology	ILAE surgical outcome	Total recording time	# seizures analysed	# electrodes analysed	Sampling frequencies	AED reduction performed?
Study 012-2	MC	37	M	B	T	Other	-	13d 16h	28	81	499.907 Hz	-
Study 017	MC	39	M	R	FT	Other	4	7d 17h	9	16	499.907 Hz	-
Study 019	MC	33	M	L	T	-	5	5d 16h	33	96	499.907 Hz	-
Study 020	MC	10	M	R	F	-	4	5d	8	55	499.907 Hz	-
Study 021	MC	16	M	R	FT	Other	1	6d 11h	13	108	500 Hz	-
Study 024	MC	23	F	B	TP, IH	-	-	8d 10h	12	83	500 Hz	-
Study 026	MC	9	M	L	F	FCD	1	3d 3h	21	81	499.907 Hz	-
Study 027	MC	34	F	L	T	HS	-	3d 21h	6	47	500 Hz	-
Study 030	MC	18	F	L	FP	FCD	3	5d 23h	8	63	500 Hz	-
Study 033	MC	3	M	L	F	TS	5	6d 17h	17	127	500 Hz	-
Study 037	MC	62	F	R	F	-	-	8d 23h	8	78	499.907 Hz	-
Study 038	MC	58	M	L	FT	-	1	3d	10	86	500 Hz	-
I002_P006_D01	HUP	26	F	R	T	-	-	12d 22h	7	83	512 Hz	yes
95	UCLH	35	M	L	OP	Other	4	7d 1h	13	56	512 Hz, 1024 Hz	no
756	UCLH	38	F	B	T	Other	3	6d 19h	6	20	1024 Hz	yes
770	UCLH	25	F	L	P	FCD	3	4d 3.6h	8	71	512 Hz	no
821	UCLH	25	F	L	T	HS, BDI	1	6d 4h	9	46	512 Hz	yes
931	UCLH	28	M	L	T	HS	4	7d	11	58	512 Hz	yes
934	UCLH	28	F	R	OP	TS	1	1d 19h	40	76	512 Hz	no
999	UCLH	28	M	L	F	FCD	1	12d 5h	26	73	512 Hz	yes
1005	UCLH	21	F	R	T	HS	2	8d 21h	15	85	512 Hz	yes
1097	UCLH	28	M	L	F	GL	1	1d 20h	8	84	512 Hz	no
1109	UCLH	31	F	R	T	CAV	1	6d 1h	13	53	1024 Hz	yes
1149	UCLH	43	F	R	TOP	DNT	1	7d 22h	24	62	512 Hz, 1024 Hz	no
1163	UCLH	27	F	L	F	FCD	1	8d	8	111	512 Hz	yes
1167	UCLH	39	M	L	P	CAV	4	5d 22h	43	51	1024 Hz	no
1168	UCLH	60	F	L	F	FCD	2	2d	10	94	512 Hz	no
1182	UCLH	28	M	R	P	FCD	3	5d 5h	52	75	512 Hz	no
1196	UCLH	41	M	R	T	HS	3	15d 22h	11	34	1024 Hz	yes
1200	UCLH	24	F	R	T	HS	1	2d 19h	14	71	512 Hz	yes
1211	UCLH	26	M	R	T	Other	3	5d 7h	20	77	512 Hz	yes

Hospitals
MC = Mayo Clinic
HUP = Hospital of the University of Pennsylvania
UCLH = University College London Hospital

Sex
M = male, F = female

Hemisphere
L = left, R = right, B = bilateral

Lobe
T = temporal, F = frontal,
P = parietal, O = occipital,
IH = interhemispheric

Pathology
FCD = Focal cortical dysplasia, BDI = Brain damage - inflammatory,
HS = Hippocampal sclerosis, TS= Tuberous sclerosis, GL = Glioma,
CAV = Cavernoma, DNT = Dysembryoplastic neuroepithelial tumour,
Other = other pathology. Dash indicates no available pathology information.

Table 2.1: **Metadata of EMU patients.** Patient identifiers are listed under "Subject." IEEG Portal patients (MC and HUP hospitals) have the same identifier as the one used by the database. Metadata was extracted from the reports provided on the IEEG Portal (MC and HUP patients) or the patient clinical reports (UCLH patients). For each patient, the following information is provided: **Hospital:** hospital at which the patient underwent presurgical monitoring. **Age:** age, in years, at the time of the presurgical monitoring. **Sex:** patient sex. **Hemisphere:** purported hemisphere of onset of the patient's seizures, based on clinical findings. **Lobe:** purported lobe of onset of the patient's seizures, based on clinical findings. Note that some patients had seizures arising from multiple lobes/at the boundary of two lobes (e.g., OP = occipital/parietal onset). **Pathology:** postoperative tissue pathology findings. **ILAE surgical outcome:** patient surgical outcome according to the International League Against Epilepsy classification (1 = seizure free, 2 = only auras, 3+ = not seizure free). A dash indicates that the patient did not undergo surgery or their surgical outcome is unavailable. For IEEG Portal patients (MC and HUP hospitals), the surgical outcome provided by the database is given. For UCLH patients, the 12 months post-surgical outcome is provided. **Total recording time:** total duration of the presurgical intracranial recording time. **# seizures analysed:** number of the patient's seizures analysed in this work. **# electrodes analysed:** number of recording electrodes included in the analysis, after removing noisy electrodes. **Sampling frequencies:** sampling frequencies at which intracranial data was acquired and stored. **AED reduction performed:** whether patient antiepileptic drug (AED)s were systematically reduced during the presurgical recording. A dash indicates that this information is unavailable.

using the seizure onsets provided on the IEEG Portal; however, seizure termination times were not marked in this dataset. Therefore, for each canine subject, the time periods around each marked seizure onset were extracted, beginning with 300s before seizure onset and ending with sufficient time after seizure onset to capture all seizure terminations, based on visual inspection (460s for canine 1, 250s for canine 2, and 150s for canine 3). Seizure termination time points were determined algorithmically after each extract was preprocessed (see Section 2.2.2). Because this identification relied on preictal data, seizures were only included in the analysis if 1) there was at least 300s between the seizure start and the termination of the previous seizure, and 2) if the preictal period, defined as three minutes to one minute before seizure start, lacked large noisy or missing segments.

Subject	Total recording time	# seizures analysed	# electrodes analysed	Sampling frequency
I004_A0001_D001 (C1)	45d 19h	86	16	399.6098 Hz
I004_A0002_D001 (C2)	451d 20h	58	16	399.6098 Hz
I004_A0003_D001 (C3)	475d 17h	43	15	399.6098 Hz

Table 2.2: **Metadata of canine subjects.** Subject identifiers are listed under "Subject." For each subject, the following information is provided: **Total recording time:** total duration of the iEEG recording. **# seizures analysed:** number of the subjects's seizures analysed in this work. **# electrodes analysed:** number of recording electrodes included in the analysis, after removing noisy electrodes. **Sampling frequencies:** sampling frequencies at which intracranial data was acquired and stored.

2.1.3 NeuroVista patients

Seizures from the NeuroVista patients were included to analyse seizure variability over longer timescales in patients with focal epilepsy. The patients and collection of their chronic iEEG data is described in detail in Cook et al. (2013). Briefly, all patients had refractory focal epilepsy and experienced 2-12 seizures per month. For the NeuroVista seizure prediction study, each patient was implanted with 16 surface iEEG electrodes over the brain quadrant thought to contain the epileptogenic zone. Additional patient details are provided in Table 2.3.

The NeuroVista seizure data from Karoly et al. (2018b) was provided for this analysis by the authors. This dataset includes iEEG of type 1 seizures (seizures with clinical manifestations and corresponding iEEG changes) and type 2 seizures (seizures with iEEG

changes comparable to type 1 seizures, but without confirmed clinical manifestations) from 12 patients. All seizure onsets occurred at least 6 minutes after the termination of the preceding seizure. We also limited our analysis to seizures that were at least 10s long, lacked noisy segments, and had sufficient data (see Sections 2.2.3 and 2.3). Patients NeuroVista 2 and NeuroVista 4 were removed from our analysis due to low numbers of analysable seizures. For the included seizures, the iEEG data from each seizure's marked onset to termination were extracted for the downstream analysis.

Subject	Age (yrs)	Sex	Age at diagnosis (yrs)	Lobe	Previous resection	# of seizures analysed	# of electrodes analysed	Total recording time (days)	Sampling frequency
NeuroVista 1	26	M	4	PT	No	94	15	767.4	400 Hz
NeuroVista 3	22	F	16	PT	Yes	273	16	557.5	400 Hz
NeuroVista 6	62	M	37	T	No	69	16	441.3	400 Hz
NeuroVista 7	52	M	26	FT	No	212	16	184.8	400 Hz
NeuroVista 8	48	M	20	FT	Yes	452	14	558.4	400 Hz
NeuroVista 9	51	F	10	OP	No	147	14	394.9	400 Hz
NeuroVista 10	50	F	15	FT	Yes	446	16	373.2	400 Hz
NeuroVista 11	53	F	15	FT	No	351	14	721.6	400 Hz
NeuroVista 13	50	M	20	T	Yes	425	16	746.9	400 Hz
NeuroVista 15	36	M	5	T	Yes	57	16	465.6	400 Hz

Sex
M = male, F = female

Lobe
T = temporal, F = frontal,
P = parietal, O = occipital,
IH = interhemispheric

Table 2.3: **Metadata of NeuroVista patients.** Patient identifiers are listed under "Subject." Clinical metadata and patient demographics are reproduced from Cook et al. (2013). The following information is provided for each patient: **Age (yrs):** patient age in years. **Sex:** patient sex. **Age at diagnosis (yrs):** patient age when they were diagnosed with epilepsy, in years. **Lobe:** purported lobe of onset of the patient's seizures, based on clinical findings. Note that some patients had seizures arising from multiple lobes (e.g., OP = occipital/parietal onset). **Previous resection**: whether the patient had undergone surgical resection prior to the chronic recording. **# seizures analysed:** number of the patient's seizures analysed in this work. **# electrodes analysed:** number of recording electrodes included in the analysis after removing noisy electrodes. **Total recording time (days):** total duration of the chronic intracranial recording time, in days. **Sampling frequency:** sampling frequency at which intracranial data was acquired and stored.

2.2 Intracranial EEG Preprocessing

Although the preprocessing pipeline was similar for all subjects, some data required additional preprocessing steps due to unique features of the data from that cohort.

2.2.1 Intracranial EEG preprocessing for epilepsy monitoring unit and canine subjects

For each subject, if different seizures were recorded at multiple sampling frequencies, all of the recordings were first downsampled to the lowest sampling frequency. Noisy channels were then removed based on visual inspection. In the remaining channels, short sections of missing values were linearly interpolated. These sections of missing values were <0.05s with the exception of one segment in seizure 2 of patient "Study 020", which was 0.514s. All channels were re-referenced to a common average reference. Each channel's time series was then bandpass filtered from 1-150 Hz (4th order, zero-phase Butterworth filter). To remove line noise, the time series were additionally notch filtered (4th order, 2 Hz width, zero-phase Butterworth filter) at 60 and 120 Hz (IEEG Portal patients and canines) or 50, 100, and 150 Hz (UCLH patients).

2.2.2 Canine subjects: Identifying seizure terminations

Because seizure end times were not marked in the canine data, seizure termination was identified algorithmically using an approach similar to Schindler et al. (2007b) after the preprocessing steps in Section 2.2.1. The seizure termination algorithm is demonstrated in an example seizure in Fig. 2.1A-C. In each channel, the time period containing seizure activity was first identified based on an increase in signal absolute slope, $S(t)$, compared to each seizure's preictal period, which was defined as three minutes to one minute before the clinically marked seizure start.

The absolute slope S of each channel i was given by

$$S_i(t) = |\tfrac{x_i(t)-x_i(t-1)}{\Delta t}|$$

where x_i is the time series voltage value of channel i and Δt is size of the time step, which was determined by the sampling frequency of the recording. $S_i(t)$ was then normalised to $S_i'(t)$ by dividing each time point by $\sigma_{i,pre}$, the standard deviation of the absolute slope of channel i during the seizure's preictal period, and smoothed by applying a 5s moving average sliding window. Channel i was considered epileptic at time point t if $S_i'(t)$ was

greater than 2.5. Seizure termination was marked as the first time, following the clinically marked seizure start, when the number of epileptic channels fell below and remained below two channels for at least 1.5s. All algorithmically marked terminations were then verified by visually inspecting the iEEG traces. We also note that because the mean and standard deviation of the absolute slope were tightly linked in our data (Fig. 2.1D), normalising the absolute slope by the standard deviation had an approximately equivalent effect to normalising by computing the z-score of the absolute slope (Fig. 2.1E).

2.2.3 Intracranial EEG preprocessing for NeuroVista patients

Seizure data was previously notch filtered at 50 Hz during the iEEG acquisition and bandpass filtered (2nd order, zero-phase Butterworth filter 1-180 Hz) by Karoly et al. (2018b). We then removed any electrodes with noisy or intermittent signal from the seizure analysis and referenced all iEEG to a common average reference.

We detected and removed periods of signal dropout by using line length to identify iEEG segments with no signal (i.e., a flat time series with no voltage changes). We defined the line length L of a time series as

$$L = \frac{1}{T-1} \sum_{i=1}^{T-1} |x_{i+1} - x_i|$$

where x_i is the ith time point in a time series with T time points.

For each seizure, time-varying line length was computed for each iEEG channel in sliding windows (1/10s window, 1/20s overlap). If a time window had at least 8 channels with line length ≤ 0.5, that time window along with the preceding and following time windows were considered missing data. Section 2.3 describes how this missing data was handled during the computation of seizure time-varying functional connectivity.

2.3 Computing Seizure Functional Connectivity

2.3.1 Choice of functional connectivity measure

To measure pairwise interactions between iEEG channels, we used a functional connectivity measure called coherence. Coherence measures the linear relationship of two signals

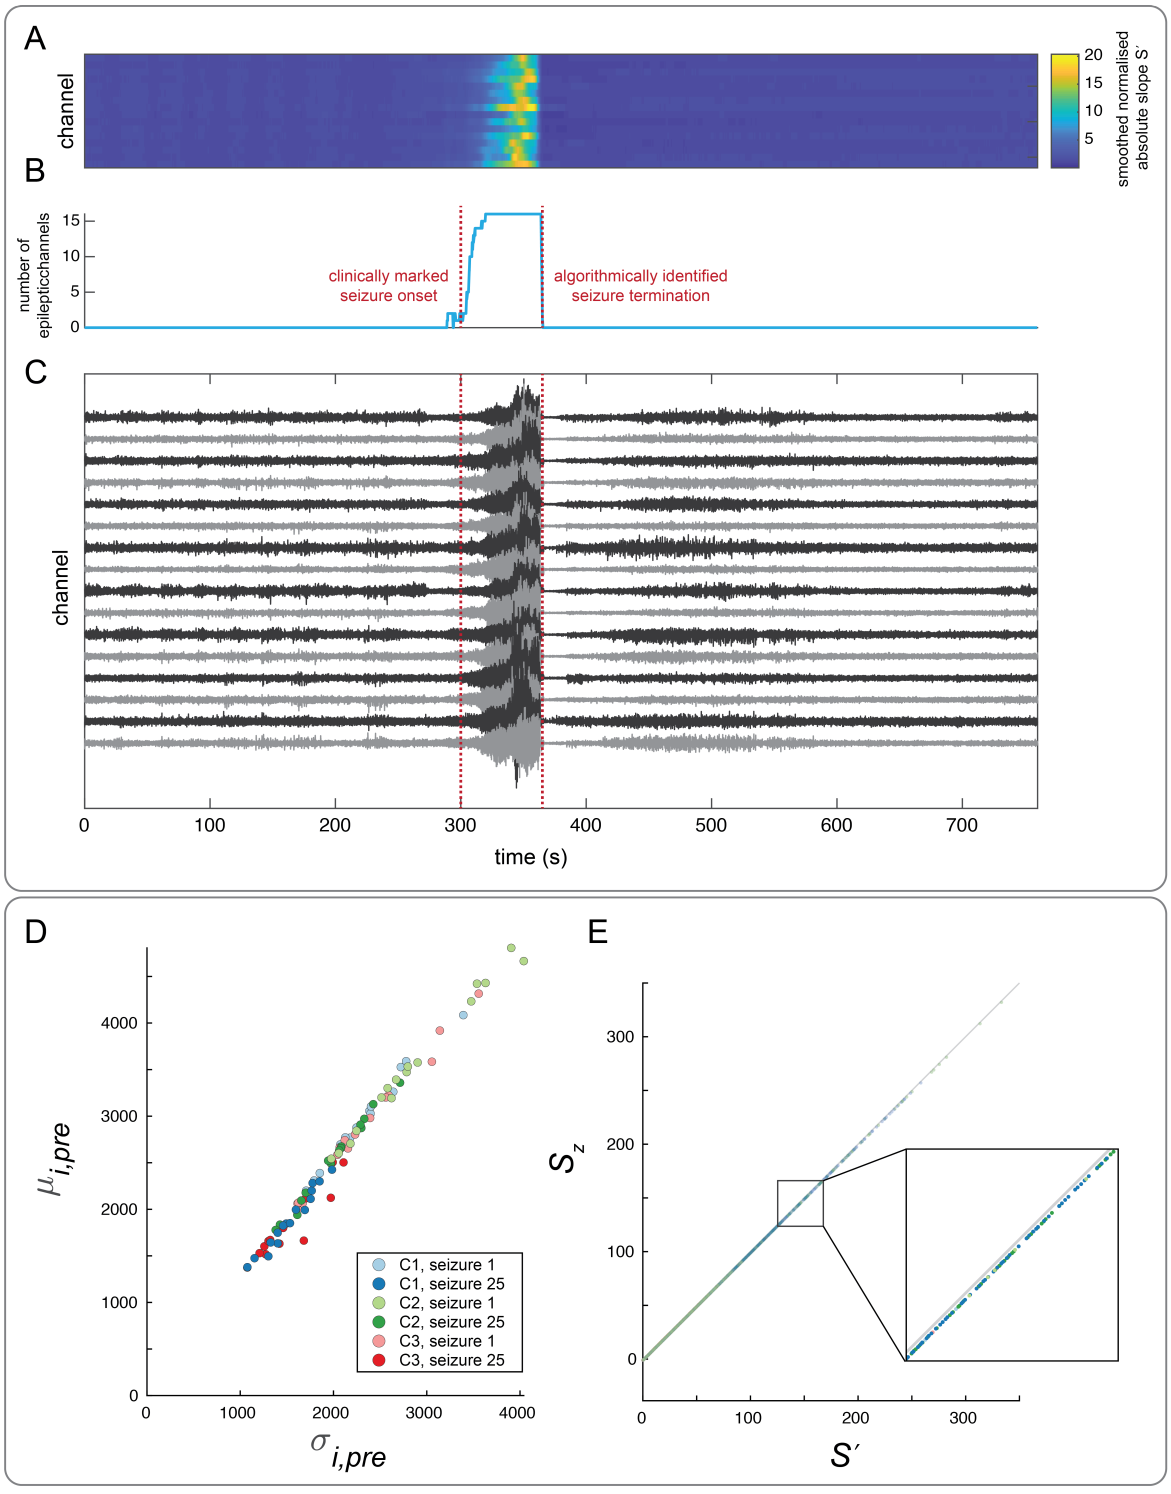
A
channel
smoothed normalised absolute slope S´
20
15
10
5
B
number of epileptic channels
15
10
5
0
clinically marked seizure onset
algorithmically identified seizure termination
C
channel
0
100
200
300
400
500
600
700
time (s)
D
$\mu_{i,pre}$
$\sigma_{i,pre}$
0
1000
2000
3000
4000
C1, seizure 1
C1, seizure 25
C2, seizure 1
C2, seizure 25
C3, seizure 1
C3, seizure 25
E
S_z
S'
0
100
200
300

Figure 2.1: **Algorithmically identifying seizure termination in canine subjects** (previous page). A-C) Identifying seizure termination in an example seizure in canine subject C1 (I004 A0001 D001). A) Heatmap of the smoothed normalised absolute slope S' of the extracted iEEG segment. B) Number of "epileptic" channels vs. time in the extracted iEEG segment. A channel was considered epileptic at a time point if its smoothed S' exceeded 2.5. The times of the clinically marked seizure onset and algorithmically identified seizure termination are marked with vertical red dashed lines. C) The iEEG of the example seizure, with the clinically marked seizure onset and algorithmically identified seizure termination marked with red lines as in (B). The algorithm successfully identified the time when high amplitude activity ceases. D-E) Comparison of normalised absolute slope S' with the z-scores of the absolute slope S_z. D) Mean μ vs. standard deviation σ of each channel i in six example preictal reference periods (two preictal periods per subject). Each point corresponds to one channel in one preictal period. E) The absolute slope z-scores S_z vs. the normalised absolute slope S', before smoothing, of the iEEG segments that corresponded to the preictal periods in (D). Each point corresponds to a channel and time point. The identity line is shown in gray. S_z and S' are closely related, with each point's S' slightly higher than its S_z; as a result, the points lie slightly below the identity line (see inset).

in the frequency domain; signals have high coherence at a given frequency if they have a consistent phase and amplitude relationship at that frequency (Bastos and Schoffelen, 2016; van Mierlo et al., 2014). We selected coherence as a measure due to this formulation in the frequency domain, its relatively low computational expense, and its frequent use as a measure of time-varying functional connectivity of iEEG signals (Burns et al., 2014; Duckrow and Spencer, 1992; Khambhati et al., 2016; Martinet et al., 2017; Shah et al., 2019; van Mierlo et al., 2014). Coherence's definition in the frequency domain allowed us to readily describe networks in different frequency bands without any additional filtering steps. This ability was an important consideration for our work, as neural interactions in different frequency bands are thought to have different functions and spatial patterns (Chapeton et al., 2019; Kopell et al., 2000). Additionally, network interactions before and during seizures depend on the analysed frequency band (Khambhati et al., 2016). Thus, we decided to analyse network interactions in several canonical frequency bands to avoid losing band-specific information about seizure network evolutions.

Other common measures for computing pairwise interactions in iEEG signals are Pearson's correlation and cross-correlation, which capture linear relationships between signals in the time domain (Bastos and Schoffelen, 2016; van Mierlo et al., 2014). Unlike correlation,

which only measures the instantaneous relationship between two signals, coherence is not decreased by phase differences between signals (Bastos and Schoffelen, 2016; van Mierlo et al., 2014). This property makes coherence well-suited for analysing brain activity, as one brain region may influence another with some time delay, resulting in phase differences between the activities of each region (Bastos and Schoffelen, 2016). As such, the most similar measure to coherence in the time domain is the cross-correlation of two signals, which allows signals to be shifted relative to each other to maximise their correlation (Bastos and Schoffelen, 2016; van Mierlo et al., 2014). Correlation and cross-correlation can also be computed in specific frequency bands by first filtering the signal using, for example, bandpass filters. However, filter frequency cutoffs are gradual, limiting the ability to isolate frequencies of interest (de Cheveigné and Nelken, 2019; Luck, 2014). Filters can also create distortions in the time domain, especially if they have steep slopes and narrow bands (de Cheveigné and Nelken, 2019; Luck, 2014). Such distortions can be particularly problematic in the presence sharp transients such as epileptic spikes (Bénar et al., 2010; Jmail et al., 2017). We therefore chose coherence as our measure due to our interest in functional connectivity in different frequency bands.

However, one disadvantage of coherence is that it is positively biased for small sample sizes; in other words, it overestimates the true relationship between two signals (Benignus, 1969). In Section 2.3.3, we demonstrate that this bias was small for the parameters we used to compute coherence. We also emphasise that there are many other approaches for measuring pairwise interactions in multivariate time series, and comparing these measures is an active area of research (Bastos and Schoffelen, 2016; Cliff et al., 2022; van Mierlo et al., 2014). Future work could determine if certain measures provide more information about the underlying pathology that contributes to seizure variability. For example, some measures may provide more information about the dynamics of the seizure onset zone (Li et al., 2021) or spatiotemporal patterns of cortical excitability (Meisel et al., 2016, 2015).

2.3.2 Computing functional connectivity

To compute the time-varying functional connectivity of each seizure, a 10s sliding window, with 9s overlap between consecutive windows, was applied to each preprocessed seizure

time series. The same sliding window parameters have previously been used to estimate time-varying coherence in seizure iEEG data (Martinet et al., 2017). For each window, the coherence between each pair of iEEG channels was computed in six different frequency bands (delta 1-4 Hz, theta 4-8 Hz, alpha 8-13 Hz, beta 13-30 Hz, gamma 30-80 Hz, high gamma 80-150 Hz). The coherence in each frequency band was computed using band-averaged coherence, defined as

$$C_{i,j} = \frac{|\sum_{f=f_1}^{f_2} P_{i,j}(f)|^2}{\sum_{f=f_1}^{f_2} P_{i,i}(f) \sum_{f=f_1}^{f_2} P_{j,j}(f)}$$

where f_1 and f_2 are the lower and upper bounds of the frequency band, $P_{i,j}(f)$ is the cross-spectrum density of channels i and j, and $P_{i,i}(f)$ and $P_{j,j}(f)$ are the autospectrum densities of channels i and j, respectively. In each window, channel auto-spectra and cross-spectra were calculated using Welch's method (2s sliding window with 1s overlap). Note that band-averaged coherence is equivalent to coherence, but filtered in the frequency domain to the frequency band of interest. As such, band-averaged coherence ranges from 0 to 1 and will be higher when the two signals have a consistent phase and amplitude relationship in the specified frequency band.

As noted in 2.2.3, many seizures in NeuroVista patients contained missing data due to signal dropouts. We tolerated some amounts of missing data in this cohort by allowing functional connectivity to be estimated using a subset of the 2s Welch subwindows (used to compute the average auto-spectrum and cross-spectra of the 10s window). Specifically, for each functional connectivity time window, we only treated the 10s functional connectivity time window as missing data if five or more of the 2s subwindows contained missing data. Any seizures with missing functional connectivity time windows were excluded from the remainder of the analysis.

Thus, in a subject with n iEEG channels, the functional connectivity of each time window was described by six symmetric, non-negative, $n \times n$ matrices, in which each entry (i, j) gives the coherence between channels i and j in the given frequency band. Each matrix was then written in vector form by re-arranging the upper-triangular, off-diagonal elements into a single column vector of length $(n^2 - n)/2$. Each vector was normalised so that the $L1$ norm (i.e., sum of all elements) was 1, thus ensuring that differences between

connectivity vectors captured a change in connectivity pattern rather than gross changes in global levels of coherence. This normalisation step also allowed the magnitude of seizure dissimilarities to be compared across subjects with different numbers of electrodes. For each time window, the six connectivity vectors were then vertically concatenated together, forming a single column vector of length $6 \times (n^2 - n)/2$. Each subject's seizure connectivity vectors were subsequently horizontally concatenated together to form a matrix V containing $6 \times (n^2 - n)/2$ features and m observations, where m is the total number of time windows across all seizures.

2.3.3 Impact of window size on bias in functional connectivity estimates

To determine the impact of bias on our coherence estimates (Benignus, 1969), we generated 100 noisy time series by randomly sampling from a standard normal distribution. A sampling frequency of 400 Hz (approximately the lowest sampling frequency observed in our cohort) was used to determine the number of samples per second of simulated data. For each unique pair of different time series (a total of 4,950 pairs), we then computed the coherence in each of the six frequency bands used in our analysis for three different time window lengths: 2s, 5s, and 10s. As described above, we used Welch's method with a 2s subwindow and 1s overlap to compute the coherence of each time window.

Fig. 2.2 shows the distributions of estimated coherence for each frequency band and time window length. Ideally, coherence estimates should be zero for pairs of noisy time series, as there is no relationship between independently and randomly sampled time series. However, we instead observed that in some scenarios, the coherence estimates could be much higher than zero. Both frequency band and time window length impacted coherence estimates; they were higher for lower frequency bands and for smaller time windows. However, at the time window length used in our analysis (10s, bottom row), the bias in coherence estimates was low, ranging from an average of 0.034 in the delta band to 0.002 for the high gamma band. Our dimensionality reduction approach (Section 2.4) also further reduced these small fluctuations in coherence due to noise.

We note that the $L1$ normalisation step will scale these small overestimates in coherence and result in different "background" levels of coherence in the final functional connectivity

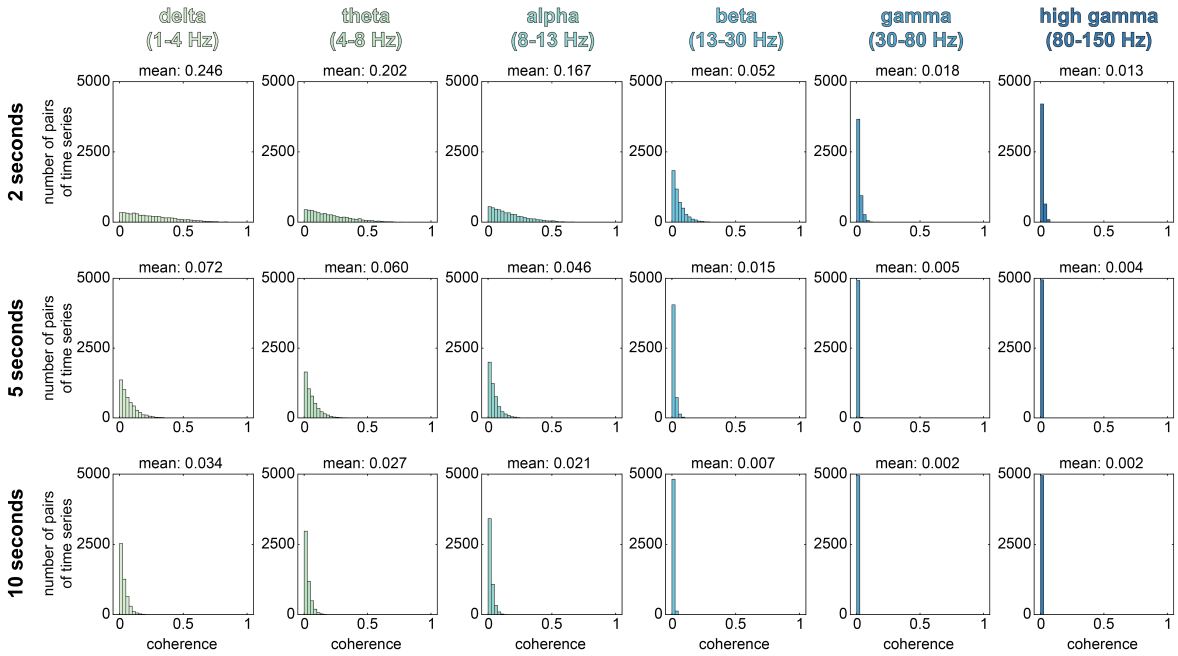

Figure 2.2: **Impact of window size and frequency band on bias in coherence estimates.** Distributions of the coherence estimates between pairs of noisy time series for different time window sizes (rows) and frequency bands (columns).

matrices. Thus, changes in background levels will also contribute to the distances between different patterns of connectivity. As such, edge-specific differences in connectivity patterns would need to be interpreted carefully. In this work, however, we do not comment on which specific functional connections vary between seizures. We instead focus on the total distance between connectivity patterns and/or differences in overall connectivity patterns (e.g., as captured by our dimensionality reduction approach, Section 2.4). If edgewise differences are evaluated in future work, it may be beneficial to first determine which connections are statistically significant (Kramer et al., 2009) to avoid analysing spurious connectivity changes.

2.3.4 Definition of seizure pathways

Throughout my thesis, I will also refer to the seizure time-varying functional connectivity as "seizure (functional) network evolutions" and "seizure pathways (through network space)." Here the term "pathway" does not refer to an anatomical pathway, but rather describes how the seizure traverses a high dimensional feature space from time window

to time window as the seizure's network dynamics change. In our work, each feature is the functional connectivity (coherence) between a pair of the patient's iEEG channels in a specific frequency band.

Due to our choice of functional connectivity measure, seizure pathways describe two aspects of seizure evolutions. First, they capture the functional connectivity edges, or the strength of relationship between different iEEG signals, during a seizure. This aspect is a spatial characteristic, as connectivity edges exist between different brain areas that each have a spatial location. Importantly, this spatial aspect does *not* capture the location of the seizure activity itself. Indeed, the connectivity of the seizure onset zone can be low during parts of a seizure (Burns et al., 2014), and the level of connectivity does not increase as the seizure spreads, but instead has a different temporal profile (Kramer et al., 2010). Thus, high values for functional network edges do not indicate participation in the seizure itself, but instead describe interactions between different brain areas during the seizure. Second, because we measure functional connectivity in multiple frequency bands, seizure pathways also depend on how brain regions are related in different frequencies. For example, the signal of two regions could be similar in the alpha, but not beta, frequency band. As above, this frequency band information does not directly describe seizure activity itself: it does not measure the frequency of seizure rhythms, but rather captures the frequencies of the functional interactions during seizures. Because seizure pathways capture both spatial and frequency aspects of the functional connectivity, differences in seizure pathways can be driven by changes in either or both aspects; for example, the functional edges could remain the same, but shift to a different frequency band.

Thus, here seizure pathways provide information about variability in seizure functional connectivity. Future studies could apply a similar approach using univariate measures that reflect the location and features of seizure activity itself, such as line length (Donos et al., 2018; Ramgopal et al., 2014), signal energy (Karoly et al., 2018b; Kramer et al., 2010), or spectral measures (Donos et al., 2018; Ramgopal et al., 2014; Wendling et al., 1996; Wu and Gotman, 1998), in order to directly compare seizure spread and rhythms.

2.4 Non-negative Matrix Factorisation

2.4.1 Properties and applications of non-negative matrix factorisation

Non-negative matrix factorisation (NMF) (Lee and Seung, 1999) is a dimensionality reduction approach that factors a two dimensional matrix of observations (e.g., seizure time windows) and features (e.g., connections between brain regions) into two matrices. One matrix contains a small number of basis vectors that each specify weights for all of the original features. The second matrix, known as the coefficients matrix, then provides the contribution of each basis vector to each of the original observations. Using these two matrices, the original data can be approximately reconstructed as a linear combination of basis vectors.

NMF differs from other matrix factorisation techniques, such as principal components analysis (PCA), due to the properties of the basis vector and coefficients matrices. In particular, NMF yields matrices that are do not contain any negative values, and each of the original observations is therefore reconstructed solely by adding together basis vectors (Lee and Seung, 1999). This property improves the interpretability of the NMF factorisation compared to approaches such as PCA that contain negative weights in both matrices. For example, a feature that is present in one PCA basis vector could be removed by substracting another basis vector, making it difficult to determine whether that feature is present in the original observation. However, unlike PCA, NMF does not produce orthogonal basis vectors unless additional constraints are added (Li and Ngom, 2013).

Depending on the data and any sparsity or orthogonality constraints, NMF basis vectors may have different properties (Mackevicius et al., 2019). First, if the rows of the coefficients matrix (with each row describing the contributions of one basis vector) are uncorrelated, each observation can be approximately described by only one basis vector. This approach essentially clusters the observations into groups defined by the dominant basis vector. The basis vectors themselves, however, can be highly correlated. In Ch. 6, we use this property of our data to cluster seizure time windows into different network states, each of which corresponds to a different basis vector. Alternatively, the basis vectors can be designed to be uncorrelated, creating non-overlapping basis vectors that

are added together to generate the original observations. In this case, a given observation cannot be described by a single basis vector, and the coefficients matrix therefore contains correlated rows. Due to NMF's non-negative constraints, these properties can be imposed using sparsity constraints (Mackevicius et al., 2019). For example, a sparse coefficients matrix will have a single high value for each observation (here, seizure time window) and very low or zero values elsewhere (e.g., see Fig. 2.3C, F), producing uncorrelated rows.

In this work, we use NMF for two purposes: to reduce noise in our data (in Ch. 3, 4, and 5) and to cluster seizure time windows into network states (in Ch. 6). The detailed methods for our applications of NMF are described below.

2.4.2 Matrix factorisation of seizure time-varying functional connectivity, V

We first used NMF to reduce noise in the connectivity matrices. In particular, small fluctuations in the functional connectivity due to noise would create a high baseline dissimilarity between seizures. Therefore, NMF was used to approximately factor each subject's seizure time-varying connectivity matrix V into two non-negative matrices, W and H, such that $V \approx W \times H$ (Fig. 2.3A-C). The matrix W contained subject-specific basis vectors, each of which had $6 \times (n^2 - n)/2$ features that captured a pattern of connectivity across all channels and frequency bands. Each original seizure time window was summarized as an additive combination of these basis vectors, with the coefficients matrix H giving the contribution of each basis vector to each time window. These factorizations were subject-specific since the basis vector features depended on the iEEG electrode layout in each subject. The optimal number of basis vectors, r, was determined using stability NMF (Wu et al., 2016).

2.4.3 Selecting the number of basis vectors using stability non-negative matrix factorisation

To determine the optimal number of basis vectors, r, for each subject, the highest r that produced consistent sets of basis vectors was found (Fig. 2.4). This approach, known as stability NMF (Wu et al., 2016), exploits the non-deterministic nature of NMF to

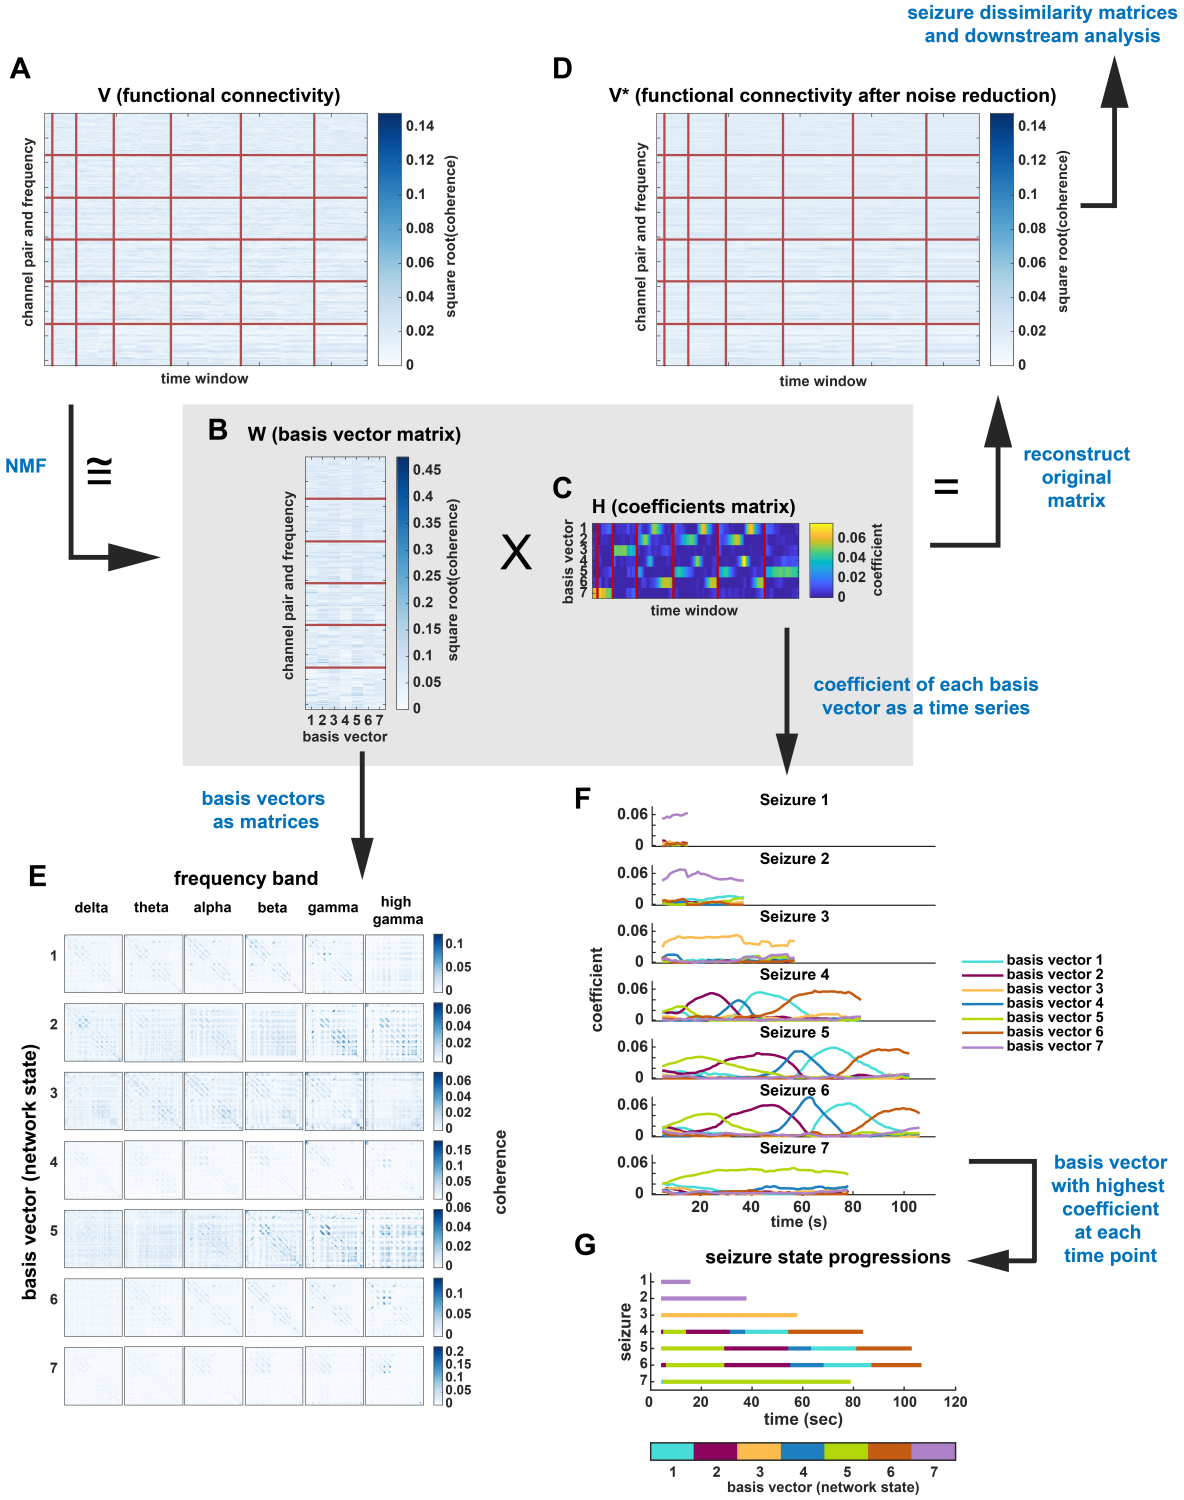
A
V (functional connectivity)
channel pair and frequency
time window
square root(coherence)
D
V* (functional connectivity after noise reduction)
seizure dissimilarity matrices and downstream analysis
NMF
≅
B
W (basis vector matrix)
basis vector
C
H (coefficients matrix)
basis vector
coefficient
X
=
reconstruct original matrix
coefficient of each basis vector as a time series
basis vectors as matrices
E
frequency band
delta
theta
alpha
beta
gamma
high gamma
basis vector (network state)
coherence
F
Seizure 1
Seizure 2
Seizure 3
Seizure 4
Seizure 5
Seizure 6
Seizure 7
coefficient
time (s)
basis vector 1
basis vector 2
basis vector 3
basis vector 4
basis vector 5
basis vector 6
basis vector 7
basis vector with highest coefficient at each time point
G
seizure state progressions
seizure
time (sec)
basis vector (network state)

Figure 2.3: **Workflow for using NMF to reconstruct seizure functional connectivity and compute progressions of seizure network states** (previous page). The factorisation of patient I002 P006 D01, which was chosen using stability NMF (see Fig. 2.4), is used as an example. From the matrix V (A), which contains the functional connectivity evolution of all seizures, NMF found a set of basis vectors, each of which forms a column in the matrix W (B), and a coefficients matrix, H (C), that describes the contribution of each basis vector to the observed connectivity at each time point. W and H were used to construct a noise-reduced version of the original matrix, V^* (D). In V, V^*, and W, rows correspond to features (here, the coherence between a given pair of channels at a given frequency band). The boundaries between features corresponding to different frequency bands are delineated using horizontal red lines (e.g., the top rows correspond to coherence in the delta frequency band). The heatmap colouring corresponds to the square root of the coherence between pairs of channels to aid in visualising the structure of the data. Meanwhile, the columns of V, V^*, and H correspond to a seizure time window, with vertical red lines marking the boundaries between seizures. E) Each NMF basis vector is a single column vector in W; however, each of the component frequency band matrices can be re-written in matrix form to emphasise the network pattern. Each basis vector therefore corresponds to a set of six connectivity matrices that describe the network configuration of all pairs of channels across the six frequency bands. Note the different colourbar scale for each basis vector. F) The coefficients matrix H can be visualised as a time series of the coefficients for each basis vector. Each time point corresponds to the functional connectivity of a 10 s window of seizure iEEG and, at each time point, the coefficients indicate how much each basis vector contributes to the observed seizure connectivity. Note that at a given time, a single NMF basis vector usually has a much higher coefficient relative to the other basis vectors (i.e., there is sparsity in each column of the coefficients matrix, H). Thus, the dominant basis vector provides a simplified description of the network dynamics at that time point. G) As an alternative visualisation of seizure evolutions, each time point was assigned to a network state corresponding to the dominant NMF basis vector, resulting in a series of state progressions for each seizure.

identify the r at which W consistently converges to a similar set of basis vectors. Since the resulting stable NMF basis vectors can be reliably found, they are thought to provide a meaningful representation of the data. To perform stability NMF for each subject, the value of r was scanned from 1 to 20. This scan range was chosen based on the observation that the stability of the factorisation greatly decreases at approximately $r > 10$ in our data, and it is consistent with the number of connectivity patterns typically found in seizure iEEG data in other studies (Burns et al., 2014; Khambhati et al., 2017, 2015). At each r, NMF of V was performed 25 times using the alternating non-negative least squares with block principal pivoting method (Kim et al., 2014; Kim and Park, 2011). Each iteration used different random initializations of W and H, thus yielding 25 different factorizations of V at each value of r. Using the method established by Wu et al. (2016), for each r, the instability I of two sets of basis vectors W and W' was defined as

$$I(r)_{W,W'} = \frac{1}{2r}(2r - \sum_{j=1}^{r} \max_{1\leq i\leq r} P_{i,j} - \sum_{i=1}^{r} \max_{1\leq j\leq r} P_{i,j})$$

where r is the number of basis vectors in both W and W', i and j are the ith and jth basis vector in W and W', respectively, P is the Pearson's cross-correlation matrix of the sets of basis vectors, and $P_{i,j}$ is the correlation between basis vectors i and j. Low values of I indicate that similar sets of basis vectors were found in the separate iterations; indeed, if the two sets of basis vectors are the same (minus reordering), then $I = 0$. The instability of all pairs of basis vector sets was then averaged to produce $I_{avg}(r)$. The highest r with $I_{avg}(r) \leq 0.005$ was selected for each subject, thus allowing small deviations between the observed basis vector sets while still enforcing consistent factorisations across iterations. At this r, the factorisation yielding the lowest reconstruction error was used to construct $V^* = W \times H$, a lower-rank approximation of the original time-varying seizure functional connectivity. This noise-reduced version of the connectivity was used in the downstream analysis.

2.4.4 Reconstruction of the seizure time-varying functional connectivity, V^*

For each subject, the selected factorization was then used to create $V^* = WH$, a lower-rank approximation of the original time-varying seizure functional connectivity (Fig. 2.3B-

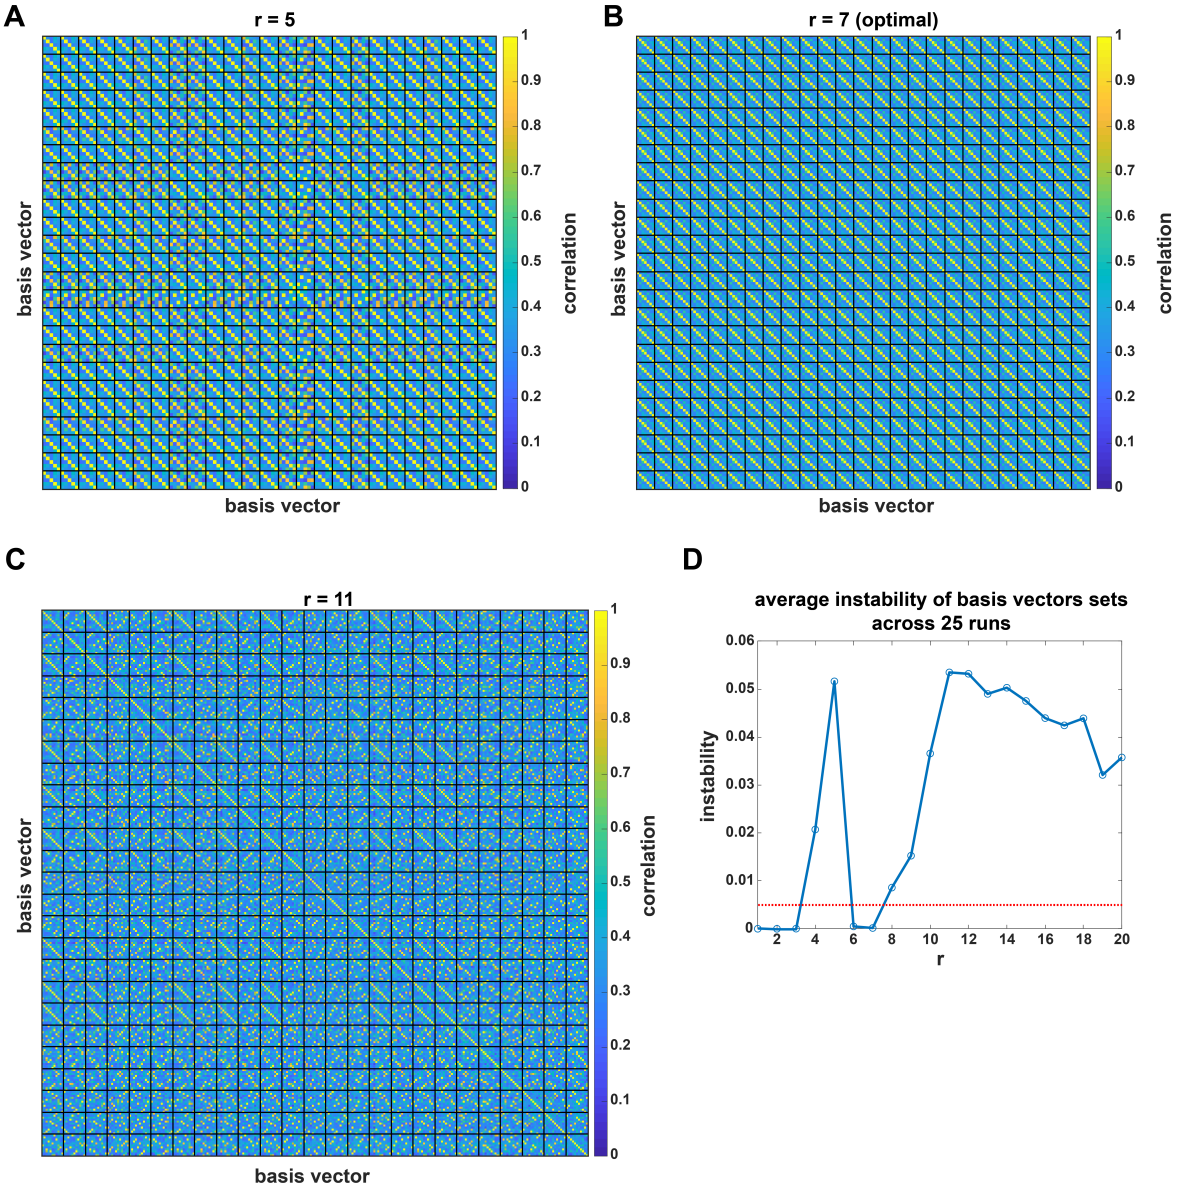
A
r = 5
basis vector
basis vector
correlation
B
r = 7 (optimal)
basis vector
basis vector
correlation
C
r = 11
basis vector
basis vector
correlation
D
average instability of basis vectors sets across 25 runs
instability
r

Figure 2.4: **Finding the optimal number of NMF basis vectors using stability NMF in an example EMU patient, I002 P006 D01** (previous page). At each number of basis vectors, NMF was repeated 25 times using different initialisations, yielding 25 sets of basis vectors. A-C) Correlation matrices showing Pearson's correlation between all pairs of basis vectors found within and between different initialisations of the NMF algorithm. In each matrix, the same number of basis vectors r were optimised for in each run: 5 (A), 7 (B), or 11 (C) basis vectors. Black lines mark the divisions between different runs of the algorithm. When possible, the basis vectors were re-ordered to emphasise similarity between different runs. Specifically, within run i, if there was a unique closest match to each of the basis vectors of run 1, the basis vectors of run i were re-ordered so that they were in the same order as the corresponding basis vectors in run 1. When $r = 7$ (B), note that similar sets of basis vectors were found across each run: in a given run, each basis vector had a high (close to 1) correlation to a basis vector of another run. Meanwhile, when $r = 5$ (A) or $r = 11$ (C), there was variability in the sets of basis vectors found across runs. D) Plot of the instability, I, of the basis vector sets, averaged across all pairs of runs, vs. the number of basis vectors, r. For a given r, the average instability quantifies the dissimilarity in the basis vectors found across runs. A low average instability indicates that similar sets of basis vectors were found regardless of the initialisation of NMF algorithm; i.e., NMF converged to similar sets of basis vectors from different initial points in the search space. We defined the optimal number of basis vectors, r, as the highest r at which $I(r) < 0.005$. For this patient, the optimal number of basis vectors was $r = 7$.

D). This return to the original feature space is necessary since NMF basis vectors are not orthogonal, and distances in NMF basis vector space are therefore not equivalent to distances in feature space. Each reconstructed connectivity vector was then renormalized to have an $L1$ norm of 1, ensuring that differences in reconstruction accuracy did not affect the distances between different seizure time windows.

2.4.5 Computing progressions of seizure states

The NMF factorisation can also be used to cluster seizure time windows into states (Brunet et al., 2004; Kim and Park, 2007) (Fig. 2.3F,G). Specifically, in our data, we observed sparse coefficients matrices, with only a single highly-expressed basis vector in a given time window. As such, the dominant basis vector provided a simplified description of the functional connectivity at that time. Therefore, in each subject, each time window was assigned to a network state corresponding to the basis vector with the highest coefficient. Each seizure could then be described as a progression of network states, enabling us to quantify how these states varied from seizure to seizure. We used this approach as an

alternative visualisation of seizure pathways on in the Zenodo supplementary for Chapters 3 and 4 (`http://dx.doi.org/10.5281/zenodo.3692923`) and to define progressions of seizure states for Chapter 6.

As noted in Section 2.4.1, the use of NMF to define seizure network states depends on sparsity in the H coefficients matrix. This property creates rows in the coefficients matrix whose values do not overlap, as different basis vectors contribute to different seizure time windows. The basis vectors in the W matrix, however, can be highly correlated, and different network states may therefore capture small differences in the network dynamics of different time windows. To demonstrate this property in our data, for each subject we found the median and maximum correlation between (1) all pairs of the subject's basis vectors in W and (2) all pairs of rows in the coefficients matrix H. For the basis vectors, a positive correlation indicates that the basis vectors represent similar connectivity patterns, a correlation close to zero reveals no relationship between the connectivity patterns, and a negative correlation indicates that connectivity in one basis vector is high when connectivity in the other is low. For rows of the coefficients matrix, a positive correlation indicates that the corresponding basis vectors have high contributions during similar time windows, a correlation close to zero indicates no relationship between the contributions of the corresponding basis vectors, and a negative correlation reveals that one basis vector has high contributions when the other does not and *vice versa* (e.g., if there are only two basis vectors that alternately contribute).

Fig. 2.5 shows the distributions of the median and maximum correlations for the W and H matrices. As expected based on our earlier visual observations (Fig. 2.3C, F), the rows of the H matrix had correlations that were negative or close to zero, reflecting that different basis vectors contributed to different seizure time windows (Fig. 2.5A,B). Meanwhile, the basis vectors of the W matrix tended to have weak to strong correlations, reflecting that many basis vectors contained overlapping patterns of functional connectivity (Fig. 2.5C,D). The highest correlation between two basis vectors was 0.804. However, the prevalence of moderate correlations revealed that most basis vectors, and thus seizure states, nonetheless described relatively different patterns of connectivity.

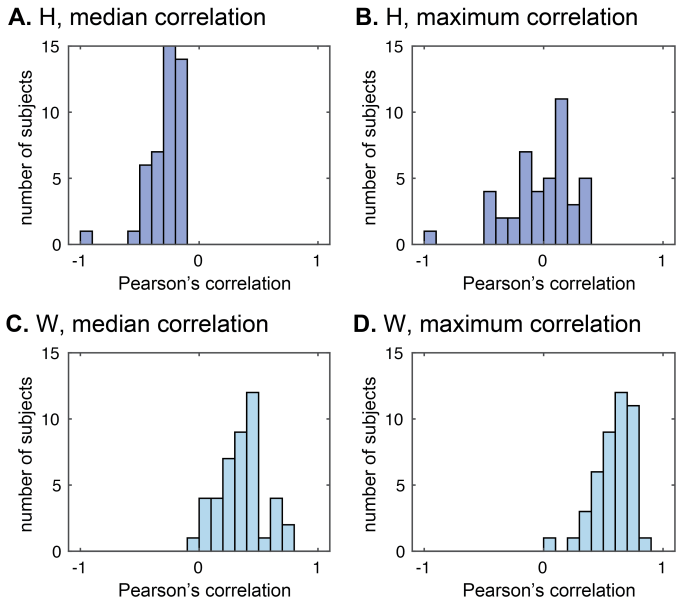

Figure 2.5: **Correlations in the W and H matrices.** A-B) Distributions of the median (A) and maximum (B) correlation between rows of the H matrix of each subject. C-D) Distributions of the median (C) and maximum (D) correlation between columns of the W matrix (i.e., basis vectors) of each subject.

2.5 Quantitatively Comparing Seizure Evolutions Using Dynamic Time Warping

2.5.1 Overview of dynamic time warping of seizure functional network evolutions

Following the NMF-based reconstruction of the seizure functional connectivity, the network evolution of each seizure was described by a multivariate time series with $6 \times (n^2 - n)/2$ features. To compare network evolutions across within-subject seizures, a seizure dissimilarity matrix was created for each subject. Each pair of seizure functional connectivity time series was first warped using dynamic time warping, which stretches each time series such that the total distance between the two time series is minimized. This step ensures that 1) similar network dynamics of the two seizures are aligned and 2) the warped seizures are the same length. We chose to minimize the $L1$ distance between each pair of seizures as this metric provides a better measure of distances in high-dimensional spaces (Aggarwal et al., 2001).

Following dynamic time warping, the $L1$ distance between the pair of warped time series was computed, resulting in a vector of distances capturing the dissimilarity in the seizures'

network structures at each time point. The seizure dissimilarity between the two seizures was defined as the average distance across all warped time points. The seizure dissimilarity matrix contains the dissimilarities between all pairs of the subject's seizures.

2.5.2 Demonstration of dynamic time warping of seizure data

In this section, we demonstrate how DTW aligns similar dynamics across seizures. In our application, this warping allows us to identify seizures with similar network evolutions, even if the rates of the evolutions differ.

DTW selectively "stretches" two time series in order to minimise the total distance between them. In our application, the distance between two time series reflects the difference between the network evolutions of two seizures. Importantly, DTW can only stretch each time series in order to align similar dynamics, which means that the algorithm cannot

1. skip time points; instead, all time points of both time series, including beginning points, must be included in the warp path.

2. repeat earlier time points; i.e., the warp path cannot double back on itself in order to repeatedly include a section of one of the time series.

The warping process is also repeated individually for each pair of seizures, and the warp length and path will therefore differ between different pairs of seizures. In order words, this algorithm provides a pairwise alignment of time series, rather than a multiple time series alignment. We focus on pairwise alignments because a global warping solution would likely sacrifice the optimal alignment of some pairs of time series.

To understand how DTW aligns pairs of time series, consider two seizures, seizure A with M windows, and seizure B with N windows. From these two seizures, we compute the time-time distance matrix D (Fig. 2.6A). This $M \times N$ matrix contains the pairwise distance between the functional connectivity of each pair of time windows across the two seizures: $D(m, n)$ is the distance between the functional connectivity of the mth time window of seizure A and the nth time window of seizure B. Thus, the row of D corresponds

to the time window index of seizure A, while the column of D corresponds to the time window index of seizure B.

DTW finds a path through this distance matrix that minimises the total distance between the two seizures. The algorithm only allows three categories of moves through the matrix (Fig. 2.6A):

1. Horizontal moves, $D(m, n) \rightarrow D(m, n + 1)$, stretch seizure A (blue arrow)
2. Vertical moves, $D(m, n) \rightarrow D(m + 1, n)$, stretch seizure B (red arrow)
3. Diagonal moves, $D(m, n) \rightarrow D(m + 1, n + 1)$, do not stretch either seizure (purple arrow)

We demonstrate this process by visualising the time-time distance matrices and warp paths of three pairs of patient 1109's seizures (Fig. 2.6B-D). Each time-time distance matrix is computed by calculating the $L1$ norm distance between the functional connectivity vectors (reconstructed following NMF) of each pair of time windows in the two seizures. Low distances reveal pairs of time windows with similar functional network dynamics. Points in the warp path (red squares), laid over the time-time distance matrix, indicate which pairs of indices were aligned. Finally, to calculate the seizure dissimilarity between two seizures, we average over the pointwise distances along the warp path. To visually demonstrate how DTW aligns similar time windows, we additionally assigned each seizure time window to a network state (section 2.4.5); each seizure can therefore be described as a progression of different network states. Time windows with the same state have similar functional network connectivity and should therefore, when possible, be aligned across different seizures.

To provide a simple example of the warping process, we first examine the warp path of seizures 7 and 12, which have a relatively low seizure dissimilarity (0.62) (Fig. 2.6B). Both seizures begin with network state 3 (orange) and then transition to network state 1 (aqua). While the duration of state 3 is the same in both seizures, state 1 lasts longer in seizure 12 than seizure 7. As such, window 7 of seizure 7 is repeated in order to align the two seizures, as shown by the horizontal line in the warp path.

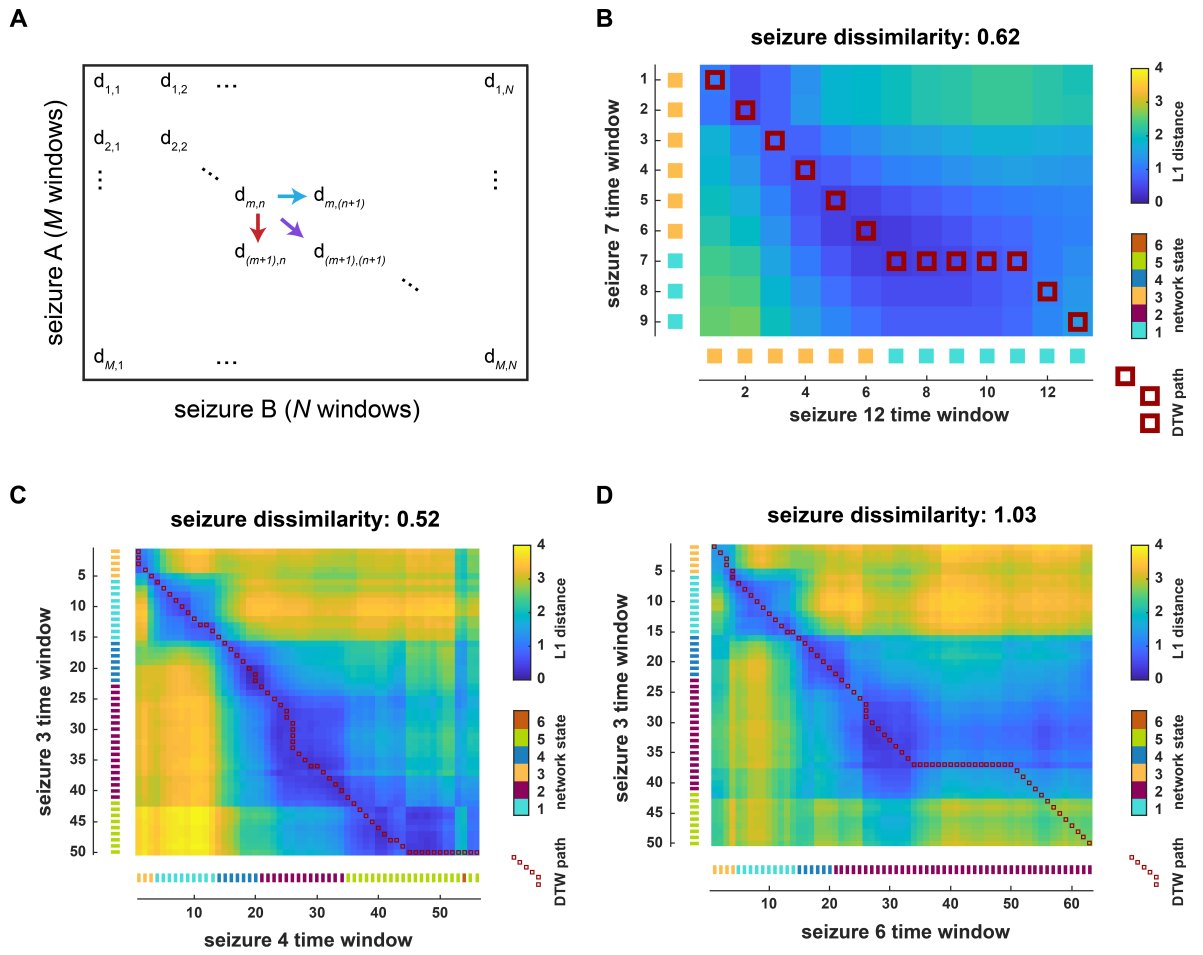

Figure 2.6: **Visualising the optimal seizure alignment found by dynamic time warping.** A) The time-time distance matrix, $D_{M\times N}$, contains the distance between each pair of time windows across two seizures. A warp path must start at $D(1,1)$ and move horizontally, vertically, or diagonally until it reaches $D(M,N)$. B-D) The warp path and time-time distance matrices of patient 1109's (B) seizures 7 and 12, (C) seizures 3 and 4, and (D) seizures 3 and 6. Each heatmap shows the $L1$ distance between the functional network evolutions of the pair of seizures. Low distances, in blue, indicate time points with similar network dynamics. Red squares indicate the warp path chosen by the dynamic time warping algorithm; if entry (i,j) is part of the warp path, then the corresponding time windows of the two seizures were aligned. For example, in (B), seizure 7's time window 7 is aligned to seizure 12's time windows 7-11 . Along each side of the time-time distance matrix, the network state of each time point is shown. Note that since the time-time distance matrix is calculated from the functional connectivity time courses, and not the simplified state descriptions, there are non-zero distances between points assigned to the same state. However, the network state progressions provide a useful visualisation for how the warp path aligns comparable network states.

Fig. 2.6C shows the warp path of seizures 3 and 4, which have similar state progressions but differ in the rate of their state progressions. However, because DTW aligns the similar parts of each seizure, there is also a relatively low seizure dissimilarity between these two seizures (0.52). These two seizures exemplify how certain windows are stretched to accommodate longer state durations in the other seizure: horizontal lines in the path correspond to places where seizure 3 is stretched, while vertical lines correspond to time points in seizure 4 that are repeated. For example, in seizure 3, a window of the last state (state 5, green) is stretched so that it matches the dynamics of state 4. Notably, the brief transition of seizure 4 to state 6 (dark orange) cannot be matched by any time point in seizure 3 because the warp path cannot skip this window or align it to earlier windows. In seizure 4, the longest warpings occur during state 3 (orange) and state 2 (burgundy red), which are both longer in seizure 3. Thus, by warping the seizure time series, our method recognises parts of the seizures that have similar network dynamics, despite differences in the rates of the seizure evolutions.

Finally, Fig. 2.6D demonstrates how DTW aligns two seizures, seizures 3 and 6, that have somewhat different state progressions. Compared to seizure 3, seizure 6 spends more time in state 2 (burgundy red) and never progresses to seizure 3's final state, state 5 (green). As such, in seizure 3, a window of state 2 is stretched. However, because the warp path must include all time points, the final part of seizure 6 must be matched to the final dynamics of seizure 3, even though there are high distances between these time points. This difference between the final seizure dynamics raises the dissimilarity between seizures 3 and 6: it is 1.03, which is almost double the previously examined dissimilarity between seizures 3 and 4.

Note that seizures 3 and 6 provide an example of when seizure pathways only partially match: the dynamics are initially very similar, and seizure 3 diverges by progressing to an additional state. In these cases, DTW cannot find a close alignment between the entire seizure progressions, and our seizure dissimilarity measure depends on the length of the mismatch relative to the overall length of the warping. In other words, had the green state lasted longer in seizure 3, the seizure dissimilarity measure would be higher in this example. In some cases, this dependence on state duration also means that the seizure

dissimilarity measure is not a metric distance.

2.5.3 Impact of functional connectivity window size and dimensionality reduction on seizure dissimilarities

To demonstrate the suitability of our functional connectivity window size and dimensionality reduction approach, Fig. 2.7 shows the impact of these choices on seizure dissimilarities in patient I002 P006 D01 (same example patient as in Fig. 2.3).

Fig. 2.7A first shows the iEEG traces of the seizures analysed in this patient. While our network approach did not explicitly capture the visually salient features of these seizures (e.g., ictal rhythms and spread), we nonetheless expected seizures with similar iEEG to also have similar underlying network evolutions and thus low, near-zero seizure dissimilarities. For example, we expected seizures 1 and 2 to have low seizure dissimilarity due to their highly similar iEEG traces. Seizure 3 also shared some similarities with these seizures 1 and 2, especially in the bottom two channels with the most prominent activity; however, the remaining channels had different patterns of activity, suggesting this seizure may differ from the earlier two. Meanwhile, seizures 4, 5, and 6 were also highly similar to each other; they were all secondarily generalised seizures with much greater spread and similar propagation patterns. Finally, seizure 7 began similarly to seizures 4-6, but did not secondarily generalise.

Fig. 2.7B demonstrates that our approach to computing seizure functional connectivity (10s sliding window, followed by NMF for dimensionality reduction) captured these visual impressions of the seizure similarities and differences. Smaller window sizes, such as a 2s sliding window, resulted in high dissimilarities between all pairs of seizures. There is some suggestion of similarities between the secondarily generalised seizures, but their seizure dissimilarities were nonetheless comparable to those of highly different seizure pairs. Increasing the window size to 5s and then to 10s gradually improved the distinction between the seizure dissimilarities of different seizure pairs and began to highlight the similarities between seizures 1 and 2 and seizures 4-6. However, the dimensionality reduction step using NMF was needed to produce the expected near-zero seizure dissimilarities within these groups of seizures. Additionally, this approach further highlighted the partial sim-

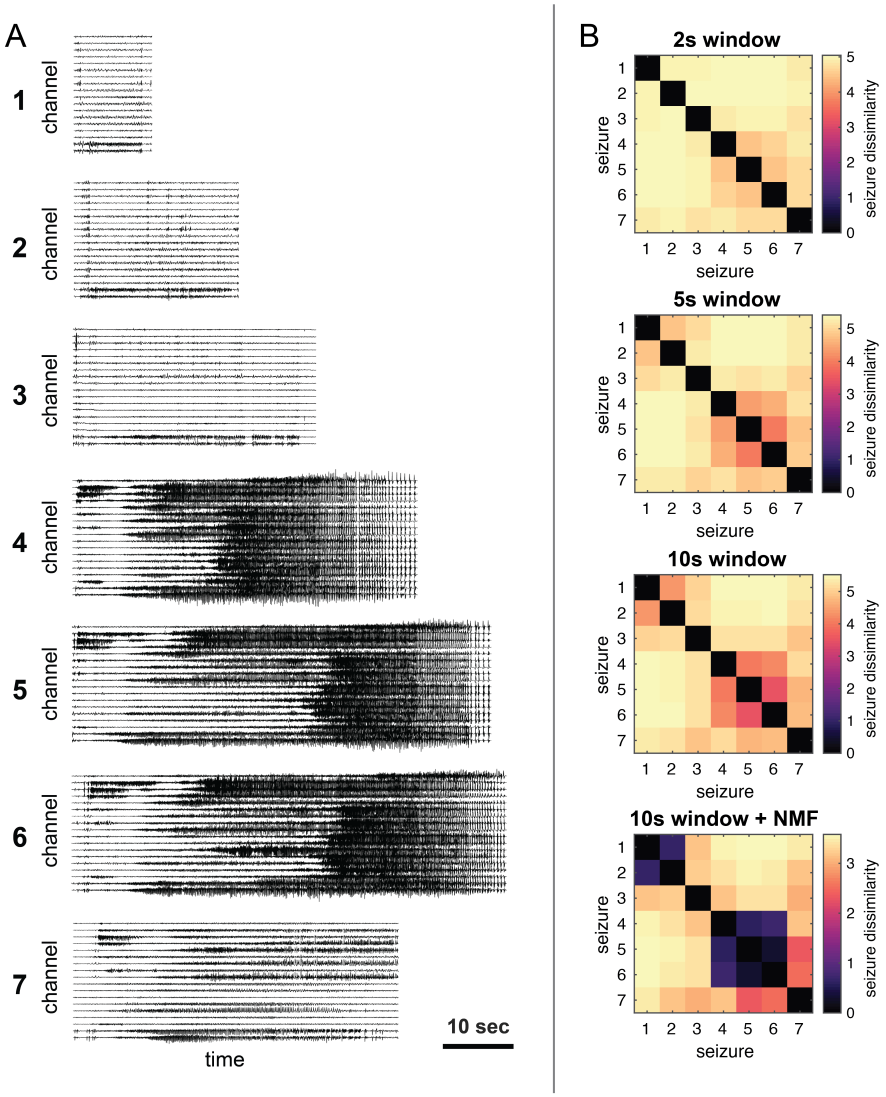

Figure 2.7: **Impact of functional connectivity window size and NMF on seizure dissimilarities in an example patient, I002 P006 D01.** A) Intracranial EEG traces of the seizures analysed for patient I002 P006 D01. For clarity, a representative subset of the channels are shown. B) Patient I002 P006 D01's seizure dissimilarities computed using four different approaches (top to bottom): 2s functional connectivity window with 1s overlap between consecutive windows, 5s window with 4s overlap, 10s window with 9s overlap, and 10s window with 9s overlap followed by dimensionality reduction using NMF. The last approach is the one used in this work. All other methods remained the same for the different approaches. Note that colourbar scales are different for each matrix to focus on differences in seizure dissimilarities within each approach, rather than the absolute values across different approaches.

ilarities between seizure 7 and seizures 4-6, resulting in an intermediate level of seizure dissimilarity.

The larger window size and dimensionality reduction step yielded the expected pattern of seizure dissimilarities by (1) improving our estimates of the functional connectivity (e.g., see Fig. 2.2) and (2) focusing on the prominent, recurring patterns of functional connectivity across seizures. In the other approaches shown in Fig. 2.7B, small differences in functional connectivity patterns due to noise or minor biological differences had large effects on seizure dissimilarities when summed across a high number of features. Our chosen approach thus minimised the influence of noise and highlighted commonalities across pairs of seizures.

2.6 Visualising Seizure Evolutions Using Multidimensional Scaling

We also wanted a way to visually compare the continuous, time-varying seizure evolutions and qualitatively assess similarities in seizure evolutions. Therefore, we used MDS to project each subject's seizure functional connectivity into a two-dimensional space such that the $L1$ distances between all pairs of seizure time windows were approximated in the projection (Sammon, 1969). Unlike matrix factorisation techniques, MDS does not transform the data in basis vectors and coefficients. Instead, variants of MDS (with the exception of Torgerson's MDS) use optimisation procedures to minimise a loss function that captures differences between the distances in the original data and distances in the MDS projection (van der Maaten et al., 2009). We used a variation of MDS called Sammon mapping (Sammon, 1969) that improves the accuracy of relatively small distances in the projection (van der Maaten et al., 2009). Sammon mapping minimises Sammon's error, E (Sammon, 1969):

$$E = \frac{1}{\sum_{i<j} d^*_{ij}} \sum_{i<j} \frac{(d^*_{ij} - d_{ij})^2}{d^*_{ij}}$$

Here, i and j are seizure time windows (columns) in the seizure time-varying functional connectivity matrix V^*, d^*_{ij} is the $L1$ distance between the functional connectivity of seizure time windows i and j, and d_{ij} is the Euclidean distance between the MDS projections of seizure time windows i and j.

Importantly, the MDS approximation of the data's original distances will contain errors that are not necessarily uniformly or predictably distributed; therefore, we used the MDS projections solely for data visualisation and not for any downstream data analysis. Specifically, the MDS projections were used to visualise seizure evolutions in Chapters 3, 4, and 5. The MDS projections of all EMU patient seizure evolutions are available in the published Zenodo supplementary files for Schroeder et al. (2020) (`http://dx.doi.org/10.5281/zenodo.3692923`).

Chapter 2: Subject Data and Computation of Seizure Network Evolutions

Chapter 3. Quantitatively Comparing Seizure Pathways Reveals Seizure Variability Within Individual Patients with Focal Epilepsy

Contents

3.1 Introduction

Focal epilepsy is characterised by spontaneous, recurrent seizures that arise from localised cortical sites (Rosenow and Lüders, 2001). An unresolved question is how much seizures themselves can vary in individual patients. Past studies suggest that seizures within a single patient share common features (Kramer et al., 2010; Schevon et al., 2012; Schindler et al., 2011; Truccolo et al., 2011; Wagner et al., 2015) and evolve through a similar sequence (Burns et al., 2014), or "characteristic pathway" (Karoly et al., 2018b), of spatiotemporal neural dynamics. However, there is also evidence that various aspects of seizures can differ within the same patient. Clinically, some patients have multiple seizure onset sites that each produce their own characteristic seizure dynamics (Spencer et al., 1981), and long-term electroencephalographic (EEG) recordings suggest that a subset of patients have multiple types of seizure evolutions (Cook et al., 2016; Freestone et al., 2017; Karoly et al., 2018b; King-Stephens et al., 2015). Ictal onset patterns (Alarcon et al., 1995; Jiménez-Jiménez et al., 2015), the extent of seizure spread (Karthick et al., 2018; Marciani and Gotman, 1986; Naftulin et al., 2018), and seizure recruitment patterns (Martinet et al., 2015) can also differ in the same patient. This variability may arise from fluctuations in the underlying brain state (Badawy et al., 2009; Bazil and Walczak, 1997; Ewell et al., 2015; Gliske et al., 2018; Khambhati et al., 2016; Naftulin et al., 2018), suggesting that background neural activity affects not only seizure likelihood (Badawy et al., 2009; Karoly et al., 2017), but also seizure evolution. Crucially, a given treatment may only address a subset of a patient's seizures: for example, a single neurostimulation protocol may not control the complete repertoire of seizures (Ewell et al., 2015) and a single prediction algorithm may fail to forecast all seizures (Cook et al., 2016; Freestone et al., 2017; Takahashi et al., 2012). Consequently, seizure variability has important implications for clinical management in these patients.

To design optimal and comprehensive treatments, we therefore need to understand the prevalence and characteristics of within-patient seizure variability. Do seizure pathways vary in all patients? Can the seizures within a patient be grouped into distinct clusters based on their evolutions? Is within-patient seizure variability solely explained by

the coexistence of multiple International League Against Epilepsy (ILAE) clinical seizure types, such as focal and secondarily generalised seizures? To answer these questions, we must objectively quantify the similarity of seizure pathways. This task is challenging due to the complexity of seizure dynamics: a variety of spatiotemporal features change independently during seizure evolution. Although some studies have quantitatively compared within-patient seizures (Le Bouquin-Jeannès et al., 2002; Louis Door et al., 2007; Wendling et al., 1997, 1996, 1999; Wu and Gotman, 1998), the current gold standard remains visual inspection of ictal EEG by trained clinicians. This latter approach is time-consuming and subjective, and can miss important features, including functional network interactions, that are difficult to detect visually.

Such functional network dynamics, also known as functional connectivity patterns, describe relationships between the activity recorded by different EEG channels. Temporal changes in network dynamics play important roles in seizure initiation, propagation, and termination (Bartolomei et al., 2004; Guye et al., 2006; Khambhati et al., 2016; Kramer and Cash, 2012; Kramer et al., 2010, 2008; Rummel et al., 2013; Schindler et al., 2007a,b, 2008; Spencer, 2002; Wendling et al., 2003), in part due to dynamic changes in the connectivity of the seizure onset zone (Bettus et al., 2008; Burns et al., 2014; Khambhati et al., 2017, 2015). Past work suggests that in some patients, the brain consistently progresses through a specific sequence of finite network states during seizures; however, other patients had unexplained variability in their seizure network evolutions (Burns et al., 2014). To fully understand how functional interactions support ictal processes, we must also understand if and how multiple seizure pathways, representing different ictal network evolutions, can co-exist in an individual patient. Such diversity would reveal that the same neural regions can variably interact to produce a variety of pathological dynamics.

In this chapter, we therefore focus on quantifying and characterising within-patient variability in seizure network evolutions; however, our approach can be adapted to compare the temporal evolutions of any ictal feature of interest. We first visualised and compared the within-patient seizure network evolutions of human patients with focal epilepsy (recorded for 43-382 hrs). In each patient, we analysed all seizures with clear electrographic correlates (i.e., both subclinical and clinical seizures). In total, we quantitatively

analysed 511 seizures (average 16.5 seizures/patient), allowing us to characterise the nature of within-patient variability in these dynamics. In each patient, we found variability in seizure network evolution, revealing that within-patient seizures are not well-represented by a single characteristic pathway through network space. However, seizures can share parts or all of the same pathway, with recurring dynamical elements across seizures. We then clustered seizures based on their within-patient dissimilarities, revealing that in most patients, seizure variability is best described as a spectrum of different evolutions, rather than distinct groups of seizure pathways. Finally, we show that the observed variability in seizure pathways is not solely a reflection of multiple clinical seizure types within individual patients. Our work provides a preliminary description of within-patient seizure variability and a quantitative tool for future research on how and why seizures change within individual patients.

3.2 Results

3.2.1 Visualising and quantifying variability in within-patient seizure pathways

Our first goal was to objectively compare within-patient seizure network evolution. For each patient, we extracted the seizure intracranial EEG (iEEG) (Fig. 3.1A) and computed the sliding-window functional connectivity, defined as band-averaged coherence in six frequency bands (Fig. 3.1B). Thus, each seizure time window was described by a set of six connectivity matrices that captured interactions between iEEG channels in each frequency band. We additionally normalised the magnitude of each connectivity matrix to focus on the evolving patterns of network interactions, rather than gross changes in the global level of coherence. The set of all possible connectivity patterns created a high-dimensional space, in which each location corresponded to a specific network configuration. As such, each time window could be represented by a high-dimensional data point, and the evolution of a seizure's network dynamics formed a pathway in this high-dimensional connectivity space. By transforming seizures in this manner, we framed our comparison of seizures as a comparison of seizure pathways (i.e., trajectories) through the high-dimensional network space.

Due to the high dimensionality of this network space, it was infeasible to directly visualise seizure pathways. However, seizure pathways could be approximated in a two dimensional projection using multidimensional scaling (MDS), a dimensionality reduction technique that attempts to maintain the distances between high-dimensional data points in the lower dimensional space (Fig.3.1C). This technique has been previously used to visualise ictal and interictal network dynamics (Khambhati et al., 2017). In our case, MDS placed seizure time windows in a two-dimensional projection based on the similarity of their network configurations; each time window was represented by a single point, and points corresponding to time windows with more similar network dynamics were placed closer together. While imperfect, this approximation of the network space nonetheless provided an intuitive visualisation for comparing seizure pathways in the same patient. For example, in patient 931, the projection demonstrated that two seizures may follow approximately the same pathway (seizures 6 and 8), part of the same pathway (seizures

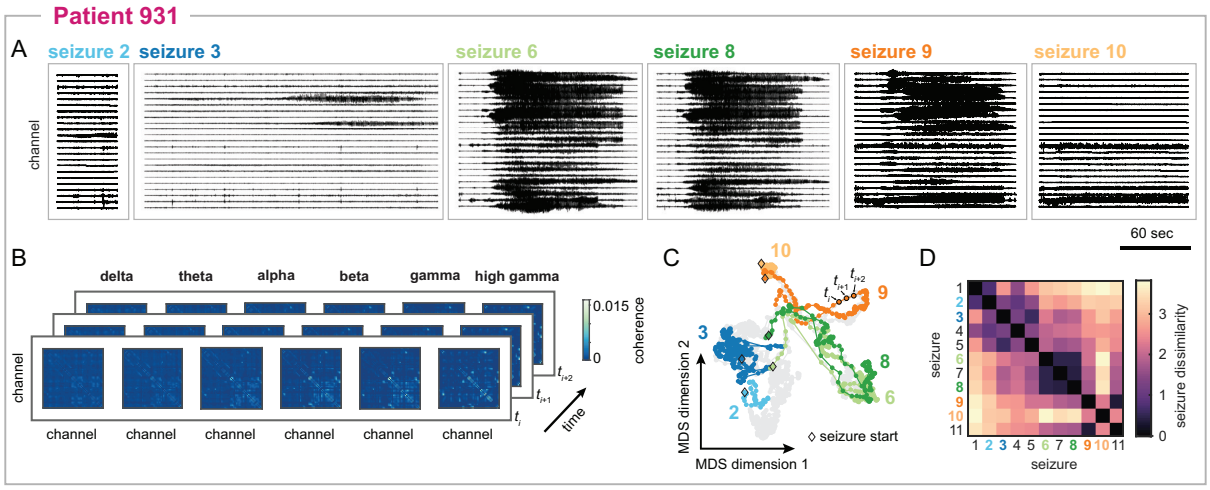

Figure 3.1: **Visualising and comparing seizure pathways through network space in an example patient, patient 931.** A) Intracranial EEG traces of a subset of the patient's seizures. For visual clarity, only a representative subset of the recording channels are shown. B) Functional connectivity of three example seizure time windows. Functional connectivity was defined as band-averaged coherence in each of six different frequency bands. Each matrix was normalised so that the upper triangular elements summed to one. Self-connections are not shown. C) Projection of all seizure time windows into a two-dimensional space using MDS, allowing visualisation of seizure pathways through network space. Each point corresponds to a seizure time window, and time windows with more similar network dynamics are placed closer together in the projection. Consecutive time windows in the same seizure are connected to visualise seizure pathways. The time windows and pathways of the six seizures shown in Fig. 3.1A have been highlighted using the corresponding colours, and the time windows of the remaining seizures are shown in grey for reference. The first time windows of the selected seizures are each marked with a black diamond. D) Seizure dissimilarity matrix of all of the patient's seizures, which quantifies the difference in the network evolutions of each pair of seizures. A low dissimilarity indicates that the two seizures have similar pathways through network space.

8 and 9), or completely distinct pathways (seizures 2 and 10) through the network space, in agreement with visual impressions of the EEG.

To quantify these visual observations, we developed a "seizure dissimilarity" measure that provided a "distance" between two seizures based on their pathways through network space. Importantly, our approach recognised similarities in seizure pathways, even if the seizures evolved at different rates, by first applying dynamic time warping (Sakoe and Seibi, 1978) to each pair of seizure functional connectivity time courses. Dynamic time warping nonlinearly stretches each time series such that similar points are aligned, thus minimizing the total distance between the two time series. We then defined the dissimilarity between two seizures as the average difference between the seizure pathways across all warped time points. The seizure dissimilarity matrix then summarised the dissimilarity between all pairs of seizure pathways in the same patient (Fig. 3.1D). In patient 931, seizures with similar pathways therefore had a low dissimilarity (e.g., seizures 6 and 8, dissimilarity 0.49); seizures with distinct, distant pathways had high dissimilarity (e.g., seizures 2 and 10, dissimilarity 3.21); and seizures with partially overlapping pathways had an intermediate level of dissimilarity (e.g., seizures 8 and 9, dissimilarity 1.75). Again, our measure of seizure dissimilarity agreed with intuitive comparisons of seizures based on visually assessing the iEEG (Fig. 3.1A) and MDS projections of the seizure pathways (Fig. 3.1C).

It is important to note that both seizure dissimilarity matrices and MDS projections were patient-specific: due to different electrode implantations, we could not compare seizures across patients using these network features. However, because we normalised the magnitude of the functional connectivity, we could compare seizure dissimilarity values across patients, even if the patients had different numbers of recording electrodes. In the remainder of the paper, we will focus on the across patients results, while using patient 931's seizures as examples. The seizure variability analysis of all patients is available on Zenodo (`http://dx.doi.org/10.5281/zenodo.3692923`).

3.2.2 Seizure variability is a common feature in all patients

Using our measure of seizure dissimilarity, we compared seizure pathways through network space in each patient. We first determined if seizure variability was present in all patients by visualising the seizure dissimilarity matrix of each patient as a distribution of seizure dissimilarities (see Fig. 3.2A for an example). Note that in these distributions, each point corresponds to the difference in network evolutions of a pair of seizures, rather than a feature of a single seizure. Fig. 3.2B shows the distribution of seizure dissimilarities in each patient, with patients sorted from lowest (patient 934) to highest (patient I002 P006 D01) median dissimilarity. Although the average level of variability differed between patients (Fig. 3.2C), it is apparent that all patients had variability in seizure network evolutions. Even in patients with more consistent seizures, such as patient 934, there were pairs of seizures with high dissimilarity, indicating dissimilar seizure pathways. Many patients, including patient 931, had varying levels of differences between pathways, with only a few pairs of similar seizures. In all patients, network differences across all frequency bands contributed to the observed seizure dissimilarities, revealing that variability in seizure network evolutions was not limited to a narrow frequency range within a given patient (Schroeder et al., 2020).

We also explored if the observed seizure variability was related to the available clinical information for each patient. We found no association between postsurgical seizure freedom (defined as ILAE surgical outcome) and median seizure dissimilarity in the 26 patients with known surgical outcomes (Spearman's $\rho = 0.06, p > 0.05$) (Schroeder et al., 2020). We also found no significant differences in median seizure dissimilarity between patients with strictly temporal ($n = 12$) and frontal ($n = 8$) onset seizures, between patients with left ($n = 15$) and right ($n = 13$) hemisphere onsets, or between male ($n = 16$) and female ($n = 15$) patients ($p > 0.05$ for all comparisons) (Schroeder et al., 2020). Sample sizes were too small to test for associations with other types of onsets (e.g., bilateral seizures) or pathological features such as hippocampal sclerosis. These findings suggest that the level of seizure variability is not associated with certain patient pathologies or treatment outcomes; instead, other factors may be more crucial for controlling the amount of seizure variability.

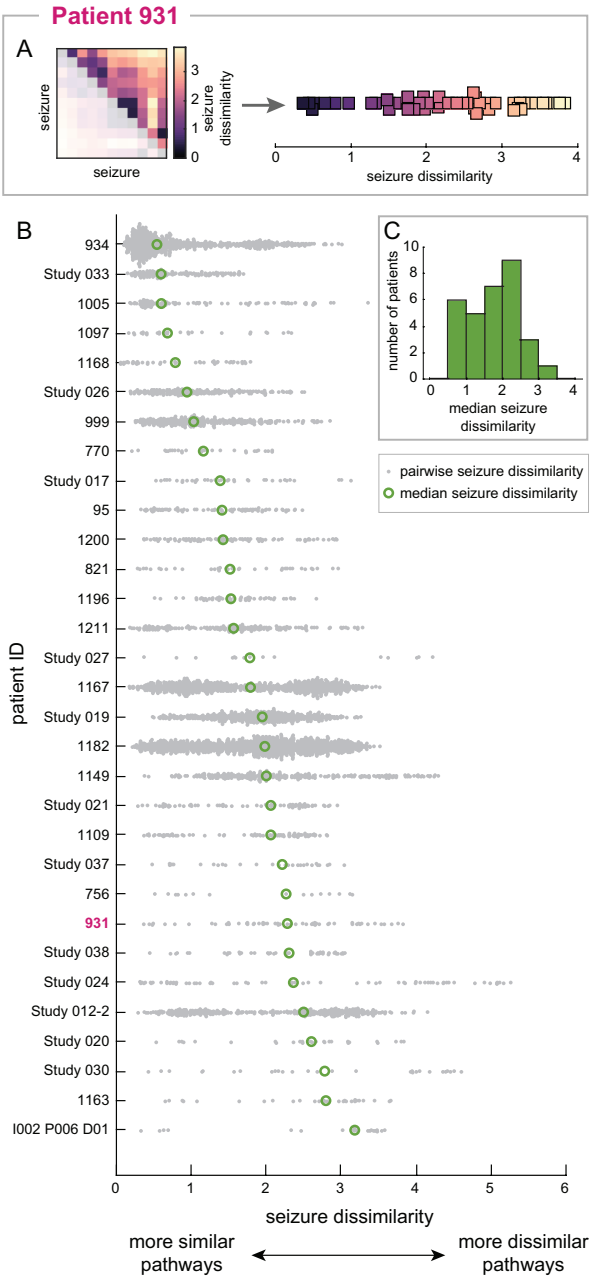

Figure 3.2: **Variability in seizure pathways is common in all patients.** A) The seizure dissimilarity matrix of patient 931 (left) was converted into a distribution of seizure dissimilarities (right) to reveal the amount and form of seizure variability in this patient. Each point in the distribution corresponds to the dissimilarity of a pair of seizures (i.e., one of the squares in the seizure dissimilarity matrix). Because the matrix is symmetric, only the upper triangular entries are plotted in the distribution. B) Distributions of seizure dissimilarities in each patient. Patients are sorted from lowest median seizure dissimilarity (patient 934) to highest median seizure dissimilarity (patient I002 P006 D01). Each grey point corresponds to the dissimilarity of a pair of seizures. The median dissimilarity of each distribution is marked by a green circle. C) Histogram of the median seizure dissimilarities of all patients.

3.2.3 Seizures form either clusters or a spectrum of different pathways within each patient

Past studies have noted that some patients have multiple populations of seizures with distinct features such as different onset sites (King-Stephens et al., 2015; Spencer et al., 1981) or durations (Cook et al., 2016; Karoly et al., 2018b). As such, we would expect the variability described in these studies to result from different, discrete seizure pathways coexisting in the same patient. However, it is also possible that seizure variability could arise from a series of small variations in seizure pathways, potentially producing a spectrum of different seizures within a given patient.

To determine which of these types of variability was present in our cohort, we found the optimal number of clusters of seizures within each patient. Fig. 3.3 shows the clustering results for each patient. In Fig. 3.3A, the seizure dissimilarity distributions from Fig. 3.2B are now coloured by whether the patient had one seizure cluster (teal) or multiple seizure clusters (purple). Patients with multiple seizure clusters tended to have more bimodal dissimilarity distributions since most of their seizure pairs had either relatively low dissimilarities (those seizure pairs that belonged to the same cluster) or relative high dissimilarities (those seizure pairs that belonged to different clusters). Importantly, patients with a single seizure cluster still had variability in seizure pathways, as they had pairs of seizures with high dissimilarities. However, without a clear way to split their seizures into different categories, the full diversity of each patient's seizure pathways could not be described by a few example seizures. In the remainder of this text and in Fig. 3.3, we will refer to these patients as having a spectrum of seizure pathways.

Fig. 3.3B shows the frequency of different numbers of seizure clusters in our cohort. Contrary to our expectation, the majority of patients (21 patients), including the example patient, patient 931, did not have distinct groupings of seizures pathways. However, as noted above, there was still variability in seizure pathways in these patients. Ten patients had two or more seizure clusters. As in the case of a single seizure cluster, there was nonetheless often additional variability in pathways within each cluster (see Supplementary section 3.5.1). The average amount of seizure variability was the same in patients with or without multiple seizure clusters (Fig. 3.3C) (two sample t-test, $p = 0.68$).

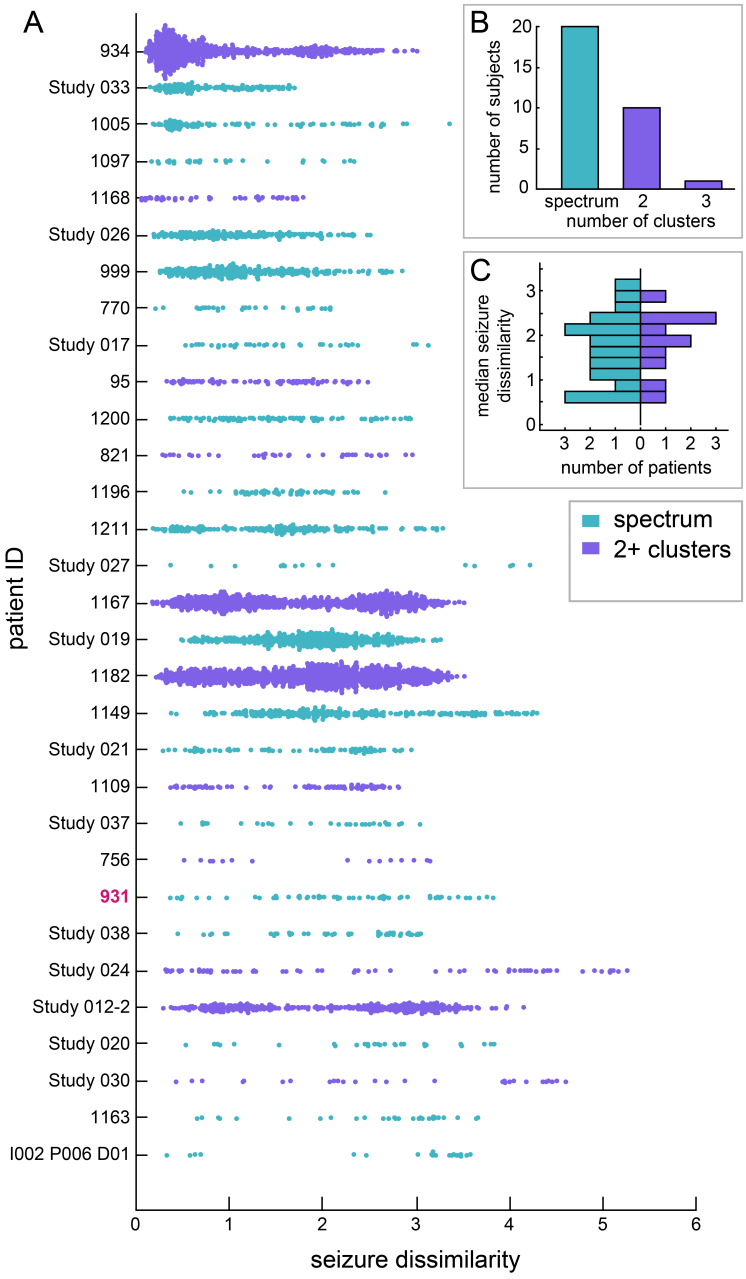

Figure 3.3: **Seizure variability may take the form of either a spectrum or clusters of different seizure pathways.** A) Distributions of seizure dissimilarities in each patient. This figure is a modified version of Fig. 3.2B. Patients are sorted from lowest to highest median seizure dissimilarity. Each point corresponds to the dissimilarity of a pair of seizures (see Fig. 3.2A for details). These distributions are now colored based on whether the seizures were best described as a spectrum of pathways (teal) or multiple clusters (i.e., separable groupings) of seizure pathways (purple). B) The number of patients with different numbers of seizure clusters based on seizure dissimilarities. C) Distribution of median seizure dissimilarities in patients with a spectrum of seizures (left, teal) or multiple seizure clusters (right, purple).

Thus, the presence or absence of different groups of seizure pathways does not indicate the average amount of seizure variability in each patient.

As in Section 3.2.2, we compared the number of seizure clusters to patient clinical features (ILAE surgical outcome, temporal vs. frontal lobe onset, left vs. right hemisphere onset, and male vs. female patients). As before, we found no significant association between this feature of seizure variability and any of these clinical variables.

To provide additional confidence in our clustering results, we repeated the seizure clustering using a different clustering method (see Supplementary section 3.5.1). The majority of patients had both the same optimal number of seizure clusters and the same seizure partitioning (i.e., assignment of seizures to clusters) regardless of the clustering algorithm. Thus, our clustering results are consistent across different types of clustering algorithms.

3.2.4 Seizure variability is not driven by differences in seizure clinical type

We next examined the relationship between variability in seizure pathways and the seizures' ILAE clinical seizure classifications (Bancaud et al., 1981; Berg et al., 2010; Fisher et al., 2017). Unlike our seizure dissimilarity measure, the ILAE clinical seizure classification was not designed to group seizures based on their specific progressions or dynamics; rather, a type is defined as "a useful grouping of seizure characteristics for purposes of communication in clinical care, teaching, and research" (Fisher et al., 2017). However, different ILAE clinical types are often characterised by specific features, such as greater spread or the presence of specific ictal rhythms (Fisher et al., 2017), suggesting that different clinical types may be distinguished by their seizure pathways. We therefore investigated the extent of amount of variability within and between different clinical types within each patient.

We labelled seizures as subclinical (i.e., electrographic only), focal, or secondarily generalised (also known as focal to bilateral tonic clonic (Fisher et al., 2017)) when that information was available in the patient's clinical report. Fig. 3.4 visualises the pathway similarity between these different ILAE clinical types in example patients. Each dendrogram shows the hierarchical clustering of a patient's seizure pathways, with seizure labels coloured by clinical type. In some patients, all recorded seizures were focal seizures (Fig.

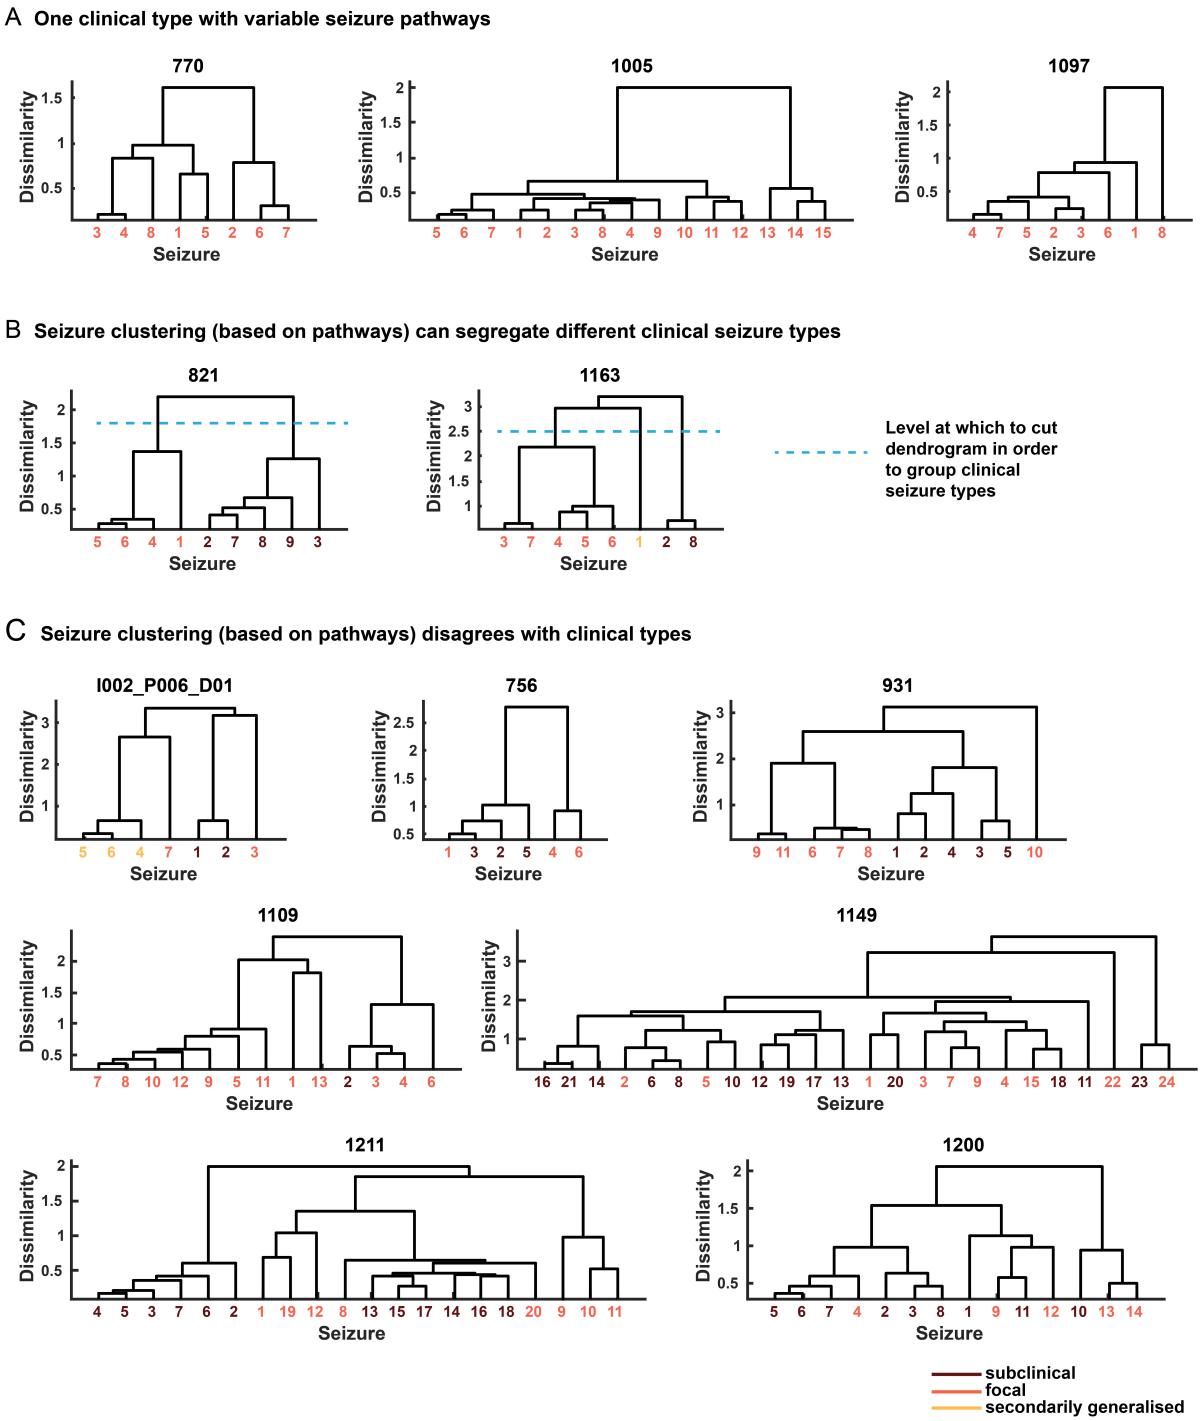
A One clinical type with variable seizure pathways
770
1005
1097
Dissimilarity
Seizure
B Seizure clustering (based on pathways) can segregate different clinical seizure types
821
1163
Level at which to cut dendrogram in order to group clinical seizure types
C Seizure clustering (based on pathways) disagrees with clinical types
I002_P006_D01
756
931
1109
1149
1211
1200
subclinical
focal
secondarily generalised

Figure 3.4: **Comparison of ILAE clinical seizure types classification and variability in seizure pathways** (previous page). A-C) Dendrograms of example patients with known clinical classifications for all seizures. Each dendrogram describes the hierarchical clustering of the patient's seizures, which was computed from seizure dissimilarities. More similar seizures, represented by leaves on the dendrogram, are joined by nodes. The height of the node linking two seizures (or groups of seizures) represents the dissimilarity between them, with higher nodes indicating less similar seizures. Seizure labels are coloured by their ILAE clinical type (dark red = subclinical, orange = focal, yellow = secondarily generalised). A) Example patients whose analysed seizures consisted of a single clinical type (focal). B) Example patients whose analysed seizures consisted of at least two clinical types, and whose hierarchical clustering of seizures agreed with the clinical seizure classification; i.e., the dendrogram can be cut at a specific level (blue dotted line) to perfectly segregate seizures of different clinical types. C) Example patients whose hierarchical clustering did not agree with the seizure clinical classification; i.e., there is no way to cut the dendrogram to perfectly segregate seizures of different clinical types. The resulting clusters will contain multiple clinical types and/or multiple clusters will contain the same clinical type.

3.4A). Although all seizures shared the same clinical type, there was nonetheless variability in the seizure pathways in each of these patients, revealing intra-patient variability within the same clinical seizure type.

In other patients, we observed that seizures clustered into groups of seizure pathways that corresponded to their clinical types (Fig. 3.4B). For example, patient 821's seizure pathways could be split into two groups, one containing only focal seizures and another containing only subclinical seizures. However, this division was not necessarily the optimal clustering of seizures that was described in Section 3.2.3, and there could still be high levels of variability within a clinical types.

Finally, in the remaining example patients, we observed that the seizures' clinical classifications did not agree with the clustering based on seizure pathways (Fig. 3.4c). In these patients, all levels of the hierarchical clustering contained clusters with multiple clinical types and/or split the same clinical type between multiple clusters. In summary, these examples demonstrate that although seizures of the same clinical type often shared similar pathways in a given patient, 1) there was often variability in seizure pathways within the same clinical type and 2) seizures of different clinical types could have relatively similar pathways, making them more similar to each other than to seizures of different clinical

types.

These findings are further confirmed by visualising the distributions of seizure dissimilarities between seizures of the same and different clinical type(s) and computing the median of each distribution (Fig. 3.5). Seizures of the same clinical type had lower median dissimilarities that seizures of different clinical types (e.g., focal vs. focal seizures, median = 1.41; focal vs. secondarily generalised seizures, median = 2.77). However, the within-type distributions nonetheless contained a wide range of dissimilarities, again revealing within-type variability in seizure pathways within individual patients. Meanwhile, pairs of seizures with different clinical types could have relatively low dissimilarities. In particular, subclincial and focal seizures often had similar pathways, resulting in a peak of low dissimilarities in their distribution. Secondarily generalised seizures, on the other hand, tended to differ from seizures of other clinical types. Therefore, while seizures of the same clinical type tended to be more similar to each other than to seizures of different types, there was also variability within clinical types and similarities across clinical types.

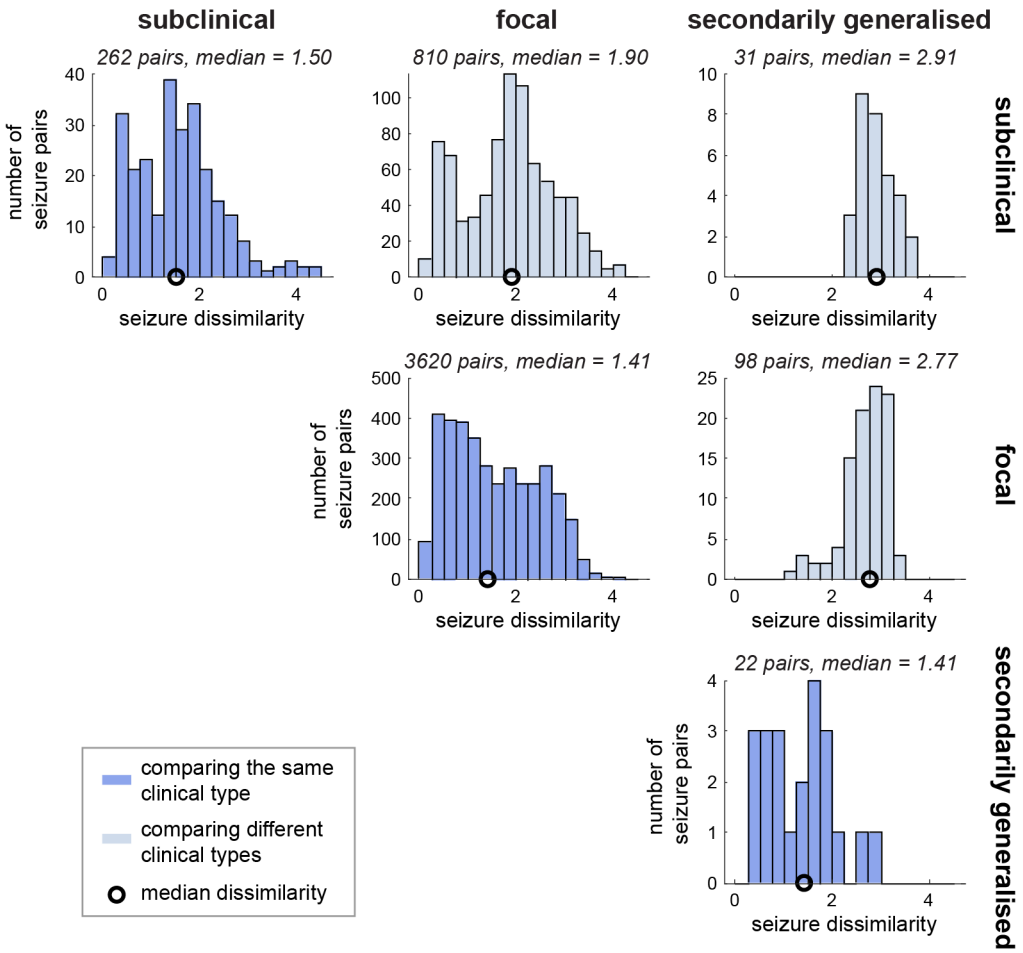

Figure 3.5: **Distributions of seizure dissimilarities between seizures of the same and different clinical types.** Each distribution contains seizure dissimilarities (computed between pairs of seizures from the same patient) of pairs of seizures that have the specified clinical type(s). Row and column labels indicate the clinical types compared in each distribution; for example, the row 1, column 1 distribution contains comparisons of subclinical seizures to other subclinical seizures, while the row 1, column 2 distribution contains comparisons of subclinical seizures to focal seizures. The number of pairs of seizures in each distribution is noted above each plot, along with the distribution's median dissimilarity (also marked with a black circle on each plot). The median dissimilarities of comparisons within seizure types (diagonal plots, in blue) are lower than the median dissimilarities of comparisons across seizure types (off-diagonal plots, in light blue); however, there is variability within and between each pair of seizure types.

3.3 Discussion

We have quantitatively compared seizure network evolutions within individual human patients with focal epilepsy, revealing that seizure variability is a common feature across patients. We often observed pairs of seizures with relatively low dissimilarity due to their largely conserved pathways through the space of possible network dynamics, suggesting that seizure evolution is not purely random. However, we likewise found that a single dynamical pathway cannot comprehensively represent all of a patient's seizure evolutions. Additionally, most patients' full spectrum of dynamics cannot be represented by a few archetypal seizures due to a lack of distinct groupings of within-patient seizure pathways. We also found that while seizures of the same ILAE clinical type tended to have more similar pathways, there is diversity within the same clinical type and commonalities across different clinical types.

We focused our investigation of seizure variability on seizure functional network evolutions due to the importance of such interactions in ictal processes (Bartolomei et al., 2004; Bettus et al., 2008; Burns et al., 2014; Guye et al., 2006; Khambhati et al., 2017, 2016, 2015; Kramer and Cash, 2012; Kramer et al., 2010, 2008; Rummel et al., 2013; Schindler et al., 2007a,b, 2008; Wendling et al., 2003) and build on previous work by demonstrating within-patient variability in these pathological network dynamics. However, in future work, the framework we present could easily be adapted to compare other features that highlight different aspects of seizure dynamics. For example, a univariate feature that captures the amplitude and frequency of ictal discharges may be better suited for comparing the involvement of different channels, similar to how clinicians visually compare EEG traces. Data from other recording modalities, such as microelectrode arrays, could be analysed to evaluate consistency in neuronal firing patterns between seizures (Schevon et al., 2012; Truccolo et al., 2011). Meanwhile, although we do not perform biophysical modelling of seizure dynamics in this work, other studies have used model inversion to hypothesise how the activities of different neuronal populations change during seizures (Freestone et al., 2014; Karoly et al., 2018b; Nevado-Holgado et al., 2012). Comparing the parameter time courses of such models could reveal different patterns of changes in neural

activity during a patient's seizures. Finally, due to patient-specific recording layouts, we focused on comparing seizure pathways within individual patients. However, comparing seizures across patients, either using spatially-independent features or common recording layouts, in future studies could uncover common classes of pathological dynamics (Jirsa et al., 2014; Karoly et al., 2018b).

To quantify within-patient variability in seizure pathways, we developed a "seizure dissimilarity" measure that addresses the challenges of comparing diverse spatiotemporal patterns across seizures. A few previous studies have attempted to quantitatively compare seizure dynamics using either univariate (Wendling et al., 1997, 1996, 1999; Wu and Gotman, 1998) or network (Le Bouquin-Jeannès et al., 2002; Louis Door et al., 2007) features computed from scalp or intracranial EEG recordings. These earlier dissimilarity measures were based on edit distance, which captures how many replacements, insertions, and deletions are required to transform one sequence into another. Importantly, unlike this previous method, our dynamic time warping approach recognises that two seizures are equivalent if they follow the same pathway, even if they do so at different rates. Despite this difference, those past studies also reported both common and disparate dynamics across within-patient seizures; however, their analysis was limited to a small number of patients and/or seizures per patient. Our work provides novel insight into the prevalence and characteristics of seizure variability by analysing over 500 seizures across thirty-one patients.

Previous work has found that within-patient seizures have similar dynamics (Burns et al., 2014; Karoly et al., 2018b; Kramer et al., 2010; Schevon et al., 2012; Schindler et al., 2011; Truccolo et al., 2011; Wagner et al., 2015), although variability may be introduced through different rates of evolution (Truccolo et al., 2011; Wenzel et al., 2017) or early termination in the seizure pathway (Karoly et al., 2018b; Wagner et al., 2015). In our cohort, we observed that subsets of within-patient seizures follow approximately the same dynamical pathway through network space, and such similar groups of seizures likely underlie these past findings. However, we also found that the complete repertoire of within-patient seizure network evolutions was poorly characterised by a single, characteristic pathway. Notably, we also found that few patients had subsets of distinct pathways; instead, small

variations between seizures often produced a spectrum of pathways. An intriguing possibility is that various "decision points," existing on the framework of potential seizure pathways, produce a repertoire of seizure evolutions. Each decision point would correspond to a seizure stage at which a seizure could evolve in multiple possible ways and/or terminate early along the potential pathway. A greater number of decision points, which in turn produce a range of small variations between seizures, would produce a spectrum of seizure pathways. On the other hand, fewer decision points and/or separate seizure pathways could produce groups of seizures with more distinct pathways. Further research is needed to map these potential seizure pathways and uncover the factors that determine how individual seizures evolve.

The crucial question is then how these different seizure pathways arise from the same neural substrate. In theory, a range of changes before or during the seizure can affect its network evolution. We hypothesise that spatiotemporal changes in the interictal neural state produce seizures with different characteristics. Past studies suggest that neural excitability (Badawy et al., 2009; Napolitano and Orriols, 2013; Wang et al., 2017), inhibition (Wenzel et al., 2017), and network interactions (Khambhati et al., 2016; Proix et al., 2018) influence certain spatiotemporal seizure features, such as the rate and extent of seizure propagation. These changes in brain state may be driven by various factors, including sleep (Bazil, 2018; Bazil and Walczak, 1997; Sinha et al., 2006), hormones (den Heijer et al., 2018; Harden and Pennell, 2013; Reddy and Rogawski, 2013; Taubøll et al., 2015), and medication (Meisel et al., 2015). Since many of these modulators vary over time, additional research is needed to determine if seizure pathways, like seizure frequency (Baud et al., 2018; Karoly et al., 2018a, 2021; Leguia et al., 2021; Maturana et al., 2020), change over different timescales within individual patients.

If interictal dynamics indeed shape how seizures manifest, future research will need to determine how specific interictal features relate to seizure characteristics. One possibility is that elements of seizure networks are activated during interictal states (Khambhati et al., 2017); thus, seizures with different network features could be preceded by preictal periods with corresponding network structures. Researchers could also relate preictal and interictal networks to other seizure features, such as seizure propagation patterns, perhaps

by investigating how the underlying structural connectome relates to functional networks (Shah et al., 2019) and mediates seizure spread (Proix et al., 2018). Importantly, the relationship between preictal network dynamics and seizure features could be limited to a specific frequency band (Khambhati et al., 2016), which could in turn suggest possible physiological mechanisms for the observed changes in seizure dynamics (Chapeton et al., 2019; Kopell et al., 2000). Additional aspects of interictal dynamics, such as the pattern of high frequency oscillations (Gliske et al., 2018) and band power changes (Naftulin et al., 2018), may also be linked to changes in seizure features. Overall, a better understanding of both functional and pathological fluctuations in interictal dynamics could suggest ways in which the background brain state alters seizure evolutions.

The clinical implications of seizure variability for patient symptoms and treatment outcomes remains an open question. We found that although seizures of the same ILAE clinical type (subclinical, focal, or secondarily generalised) tended to have more similar pathways, this classification alone misses much of the observed variability in within-patient seizures. In other words, the within-patient variability in seizure pathways in our cohort was not solely due to the coexistence of multiple clinical seizure types within the same patient. Past studies have also observed that within the same patient, seizures of the same clinical type may have different features (Fisher et al., 2017; Marciani and Gotman, 1986; Sinha et al., 2006), while seizures of different clinical types can share elements of their dynamics. For example, some patients have multiple populations of subclinical seizures, only some of which share the same onset location as the focal seizures (Farooque and Duckrow, 2014). Likewise, focal and secondarily generalised can share the same onset dynamics in the same patient (Marciani and Gotman, 1986). Since we compared within-patient seizure network evolutions, rather than propagation patterns, further work is needed to determine how commonalities in network evolutions across different clinical seizure types relate to shared seizure onset and propagation dynamics. In particular, such similarities in network evolutions could potentially reveal how the same pathological brain networks can support seizures with different clinical characteristics and produce seizures with differing spread and severity (Khambhati et al., 2016).

One limitation of our clinical analysis was that we compared our seizure pathway group-

ings to a coarse clinical classification of seizure types. A number of finer distinctions of clinical types also exist based on symptoms such as loss of consciousness, specific motor characteristics, and cognitive impairments or experiences (Fisher et al., 2017; McGonigal, 2020; Noachtar and Peters, 2009; Rossetti and Kaplan, 2010). Since seizure semiology may indicate seizure lateralisation and spread (Noachtar and Peters, 2009; Rosenow and Lüders, 2001; Rossetti and Kaplan, 2010), it is possible that variability in seizure network evolutions may more closely align with variability in seizure semiological progressions. However, since seizures with different onsets can share the same semiology (Vaugier et al., 2017), it is unlikely that there is a one-to-one correspondence between seizure pathways and seizure semiology. Although we lacked the data to pursue this analysis in our work, such analysis could help reveal the network underpinnings of different types of seizure symptoms. This knowledge could, in turn, suggest mechanisms for preventing more severe symptoms such as loss of consciousness.

Contrary to the expectation that high levels of seizure variability may worsen surgical outcomes, we found no association between these patient features. It may be that only some types of variability, such as multifocal (Ryzi et al., 2015; Spencer et al., 1981) or secondarily generalised (Baud et al., 2015) seizures, impact the likelihood of seizure freedom following surgery. Importantly, variability in the seizure onset network state does not indicate that a patient has multifocal seizures, as different network configurations can be associated with the same apparent ictal onset zone. Additionally, variability in seizure pathways may not be inherently deleterious, as long as it is observed and accounted for when planning the surgical resection. Indeed, due to the short presurgical monitoring time and limited spatial coverage of the recording electrodes, some potential seizure pathways may not have been captured (King-Stephens et al., 2015; Ung et al., 2016), leading us to underestimate the level of variability in some patients.

Although the amount of seizure variability was not associated with post-surgical seizure freedom, it may have implications for clinical treatments. First, regardless of the source of the observed seizure variability, the different seizure dynamics observed during presurgical monitoring provide crucial information for guiding surgical resection. For example, recent studies suggest that seizure network properties can help identify epileptogenic tissue

(Burns et al., 2014; Goodfellow et al., 2016; Kini et al., 2019); however, we must determine if seizures with different network evolutions provide equivalent localisation information. Seizure variability may also have implications for seizure prediction. In particular, in that same patient, seizures with different pathways may have distinct preictal signatures, making seizure prediction more difficult (Cook et al., 2016; Freestone et al., 2017). A successful seizure prediction algorithm would either need to recognise multiple signatures or find common features among the disparate preictal dynamics. Finally, neurostimulation offers a promising new approach for controlling seizures; however, in rodent models, the effectiveness of a given stimulation protocol depends on the preictal brain state (Ewell et al., 2015). Thus, such interventions may need to recognise and adapt to the specific characteristics of each corresponding seizure evolution in order to control all seizures. Importantly, our cohort was limited to patients with medication refractory focal epilepsy who were candidates for surgical resection. The characteristics and clinical implications of seizure variability may be different in other patient cohorts.

In summary, we have developed a method for quantitatively comparing seizure pathways within individual patients. Using this approach, we have shown that there is within-patient variation in seizure network evolutions in patients with focal epilepsy. In most patients, this variability creates a spectrum of possible seizure pathways, suggesting that there are many possible ways that pathways can vary within a patient. Additionally, our measure of seizure dissimilarities provides more information about seizure dynamics than high-level clinical seizure groupings. Further research is needed to determine whether and how preictal dynamics influence seizure pathways. Uncovering these mechanisms could provide novel approaches for predicting and controlling seizures that are tailored to the complete repertoire of pathological neural dynamics in each patient.

3.4 Methods

Full descriptions of the analysed patients and seizure preprocessing, as well as the computation, visualisation, and comparison of seizure functional connectivity, is provided in Chapter 2. A brief summary of these approaches and the detailed chapter-specific methods are provided below.

3.4.1 Computing seizure functional connectivity

This work analysed 31 patients with refractory focal epilepsy whose iEEG was acquired during presurgical evaluation at epilepsy monitoring units (EMUs). Briefly, for each patient, seizures were extracted and preprocessed by re-referencing the seizure iEEG to a common average reference, bandpass filtering from 1-150 Hz, and notch filtering at either 60 and 120 Hz (IEEG Portal patients) or 50, 100, and 150 Hz (University College London Hospital (UCLH) patients). The time-varying functional connectivity, defined as coherence in six frequency bands (delta 1-4 Hz, theta 4-8 Hz, alpha 8-13 Hz, beta 13-30 Hz, gamma 30-80 Hz, high gamma 80-150 Hz), was computed for each seizure using a 10s sliding window with 9s overlap between consecutive windows. Each coherence matrix was re-expressed in vector form by re-arranging the upper-triangular, off-diagonal elements, and the vector was normalised to have an $L1$ norm equal to 1. Non-negative matrix factorisation (NMF) (Lee and Seung, 1999) was then applied to each patient's set of seizure coherence matrices to reduce noise. All downstream analysis was performed on the resulting reconstructed, noise-reduced time series.

3.4.2 Visualising seizure functional connectivity

To visualise how connectivity changed over time within and between seizures, the seizure time-evolving connectivity patterns were projected into two-dimensional space using Sammon mapping, a variation of MDS (Sammon, 1969). The mapping approximated the $L1$ distances between the functional connectivity patterns of each pair of seizure time windows.

3.4.3 Computing seizure dissimilarities

For each patient, dynamic time warping (DTW) was applied to each pair of seizure functional connectivity time series such that the total $L1$ distance between the pair of time series was minimised. Their "seizure dissimilarity" was defined as the average distance between the warped time series. The dissimilarities between all pairs of seizures in a patient formed the patient's seizure dissimilarity matrix.

3.4.4 Seizure clustering and cluster evaluation

To identify groups of similar seizures in each patient, each patient's seizures were hierarchically clustered by using the seizure dissimilarity matrix as input for an agglomerative hierarchical clustering algorithm, unweighted pair group method with arithmetic mean (UPGMA). The hierarchical clustering resulted in a dendrogram that summarised the similarity between the patient's seizures. Note that the hierarchical clustering representation was an approximation of the seizure dissimilarities that forced all dissimilarities into a metric space.

The gap statistic (Tibshirani et al., 2001), which compares the within-cluster dispersion of a given clustering relative to a reference (null) distribution, was then used to determine if optimal flat (i.e., non-hierarchical) clusters of seizures existed in each patient. In order to generate reference datasets, the patient's seizures were first projected into Euclidean space using classical (Torgerson's) MDS (Fulcher, 2017). This step differs from the earlier visualisation of seizure pathways, which projected seizure time points, rather than seizures themselves. Given the seizure dissimilarity matrix, MDS assigned a coordinate point to each seizure while attempting to preserve the specified dissimilarities between seizures. In order to most closely approximate the dissimilarities matrix, the seizures were projected onto the maximum possible number of dimensions; note, however, that like the hierarchical clustering, MDS also provided a metric approximation of the nonmetric dissimilarities. One thousand reference datasets were then generated by drawing coordinates from a uniform distribution placed over a box aligned with the principal components of the projected seizure data. Each reference dataset was hierarchically clustered by computing

the distances between the coordinate points and applying the UPGMA algorithm. To test for flat clusters in the seizure data and reference datasets, the dendrograms were cut at different levels to generate $1, 2, \ldots . s$ clusters, where s is the number of seizures. At each number of clusters k, the gap statistic $G(k)$ was computed by comparing the within-cluster dispersion of the observed seizures and the reference datasets. The multiple reference datasets also allowed calculation of the standard error of the gap statistic at each k, $SE(k)$. The optimal number of clusters was defined as the smallest number of clusters where $G(k) \geq G(k+1) - SE(k+1)$, which identifies the point at which increasing the number of clusters provides little improvement in the clustering of the data (Tibshirani et al., 2001).

3.5 Supplementary

3.5.1 Supplementary clustering results

Amount of variability within and between seizure clusters

As described in the main text, each patient either had a spectrum of seizures (i.e., one seizure cluster, with variability in seizure pathways within that cluster) or multiple seizure clusters (i.e., their seizure pathways could be grouped into different classes of dynamics, with more similarity within a group than between groups) (Section 3.2.3). Fig. 3.6 shows the median seizure dissimilarity of pairs of seizures within the same cluster and the median seizure dissimilarity of pairs of seizures from different clusters in both patients with a spectrum of seizures (teal histogram) and patients with multiple seizure clusters (purple histograms).

In patients with a single seizure cluster, the median within-cluster seizure dissimilarity is equivalent to the patients' overall median seizure dissimilarity. As we saw before in Fig. 3.3C, despite having the same clustering results, the average amount of variability in seizure pathways varies in these patients.

Overall, in patients with multiple seizure clusters, the average between-cluster seizure dissimilarity is higher than the average within-cluster seizure dissimilarity, as expected. However, there is overlap in the distributions of within- and between-cluster median seizure dissimilarities. As such, depending on the patient, seizures in the same cluster can be relatively different, while seizures in different clusters can be relatively similar. Thus, a single cluster may represent seizures that may still have variable pathways, but that are nonetheless more similar to each other than to the patient's other seizure pathways. Meanwhile, in some patients, different groups of seizure pathways (corresponding to different clusters), are still relatively similar, indicating overall lower variability in seizure evolutions.

Clustering using alternative algorithm (k-means)

To provide additional confidence in our clustering results, we repeated the seizure clustering analysis from the main text using k-means clustering instead of hierarchical clustering.

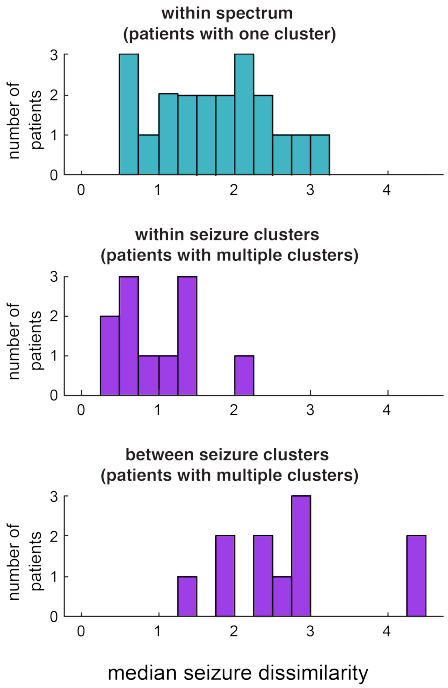

Figure 3.6: **Amount of variability within and between seizure clusters.** Teal correspond to data from patients with a spectrum of seizures (i.e., single seizure cluster), while purple corresponds to data from patients with two or more seizure clusters. Top histogram: median within-cluster seizure dissimilarity in patients with a spectrum of seizures (one seizure cluster). Middle histogram: median within-cluster seizure dissimilarity in patients with multiple seizures clusters. Bottom histogram: median between-cluster seizure dissimilarity in patients with multiple seizure clusters.

Because the k-means algorithm requires the observations to be in a metric distance space, we performed this clustering on the MDS embedding of the seizures (see "Seizure clustering and cluster evaluation"), rather than using the seizure dissimilarity matrix. As in the main text, the gap statistic was then used to determine the optimal number of seizure clusters. Of the 31 patients, 28 patients had both the same optimal number of seizure clusters and the same seizure partitioning (i.e., assignment of seizures to clusters) regardless of the clustering algorithm (hierarchical clustering or k-means clustering). Thus, in the majority of patients, we obtained equivalent clustering results using a different clustering algorithm

We next looked at the three patients whose clustering solutions differed. The two patients whose optimal number of clusters changed were patients 1167 and 1182, who had two clusters based on hierarchical clustering, but four and eight clusters, respectively, using k-means clustering. We compared these finer partitions to the hierarchical clustering of each patient by cutting each patient's dendrogram at the level required to make four and eight clusters, respectively. The different clustering solutions can then be compared using the Rand index, which is the proportion of pairs of elements (here, seizures) that have the same relative clustering assignment in both solutions; i.e., the two elements are in the same cluster in both solutions or are in different clusters in both solutions. Additionally, the adjusted Rand index (ARI) can also be computed to determine if the level of agreement is likely to arise by chance. The Rand index and ARI will both be one if the two clustering solutions are exactly the same. Meanwhile, an ARI close to zero indicates chance levels of agreement between the clustering solutions. Although these clustering solutions differed from the k-means clusters, there was a high level of overlap (Rand index of 0.78 for patient 1167 and 0.91 for patient 1182), indicating that both clustering solutions nonetheless capture similar information. The ARI (0.57 for patient 1167 and 0.69 for patient 1182) revealed that these similarities were higher than expected by chance. Finally, Patient Study 012-2 had two optimal clusters in both cases, but the partitions slightly differed; namely, the cluster assignment of two of the patient's 28 seizures changed, resulting in a Rand index of 0.86 and ARI of 0.72.

Chapter 4. Seizure Pathways Change over Circadian and Slower Timescales in Patients with Focal Epilepsy

Contents

4.1 Introduction

In patients with focal epilepsy, the frequencies of pathological neural dynamics such as seizures and interictal spikes change over time within individual patients (Karoly et al., 2016). These rates do not randomly fluctuate; rather, in many patients, seizure and spike rates change over circadian and/or multi-day cycles (Baud et al., 2018; Karoly et al., 2018a, 2021; Leguia et al., 2021; Maturana et al., 2020) that are thought to be driven by time-varying factors (Rao et al., 2020) such as sleep (Bazil, 2018; Bazil and Walczak, 1997; Sinha et al., 2006) and hormones (den Heijer et al., 2018; Harden and Pennell, 2013; Reddy and Rogawski, 2013; Taubøll et al., 2015). Importantly, understanding patient-specific cycles of seizure risk improves seizure forecasting (Karoly et al., 2017; Proix et al., 2021) and could provide opportunities for delivering time-sensitive treatments for controlling seizures (Baud and Rao, 2018; Navis and Harden, 2016; Ramgopal et al., 2013).

In the previous chapter, we demonstrated that seizure pathways, defined as seizure functional network evolutions, also vary within individual patients. However, it is unknown whether these different pathways occur randomly or if seizure pathways change over the same timescales that influence seizure occurrence and spike rates. Past studies suggest that at least some seizure features may fluctuate over different timescales. For example, in some patients, seizures are more likely to secondarily generalise during sleep (Bazil and Walczak, 1997). Additionally, some studies have found shifts in seizure onset and/or propagation dynamics in intracranial EEG (iEEG) recordings in humans (Gliske et al., 2018) and canines (Ung et al., 2016). Importantly, such changes in seizure dynamics could potentially impact seizure prediction (Cook et al., 2016) and responses to treatments such as neurostimulation (Ewell et al., 2015). Despite these observations, a more comprehensive analysis of temporal changes in seizure dynamics is lacking.

We addressed this knowledge gap by exploring how different seizure pathways were distributed in time within individual patients with focal epilepsy. In Chapter 3, we developed a method for quantitatively comparing seizure pathways within patients using a "seizure dissimilarity" measure. In this chapter, we used the previously computed seizure dissimilarities of 31 epilepsy monitoring unit (EMU) patients to determine how seizure pathways

change during short-term presurgical iEEG recordings. We first asked if seizures with more similar pathways, as described by seizure dissimilarities, tended to occur closer together in time. We then explored how seizure pathways change over different timescales and determined if patterns of variability could be explained by changes over circadian and/or slower timescales.

4.2 Results

We analysed temporal changes in seizure pathways during presurgical iEEG recordings in 31 patients with focal epilepsy. This analysis is a continuation of the comparison of seizure pathways from Chapter 3 and uses the same patient data and approaches for comparing seizure pathways. As before, we defined seizure pathways as seizure trajectories through functional network space and quantified the "dissimilarity" between each pair of a patient's seizures using dynamic time warping (DTW). This objective comparison of within-patient seizure pathways allowed us to quantify patterns of changes in seizures in the following sections.

4.2.1 Seizures with more similar pathways tend to occur closer together in time

We first asked whether seizures that occurred closer together tended to have more similar dynamics in each patient. Fig. 4.1A shows the pathways of patient 931's seizures, as well as the time that each seizure occurred relative to the patient's first seizure. From this visualisation, we see that the pathways gradually migrated through network space as the recording progressed, creating the observed spectrum of network evolutions. Moreover, looking at the seizure timings, we also see that seizures with similar pathways, such as seizures 6-8, tended to occur close together in time.

To quantify this temporal relationship, we defined a "temporal distance" matrix as the amount of time elapsed between each pair of the patient's seizures (Fig. 4.1B). Patient 931's seizure dissimilarity and temporal distance matrices have strikingly similar structures: groups of seizures with low dissimilarity tended to occur together in a relatively short time interval. In this patient, there was a strong and statistically significant positive correlation between these features (Spearman's $\rho = 0.69$, $p = 0.001$, one-tailed Mantel test), indicating that seizures with more similar pathways tended to occur closer together in time.

Fig. 4.1C summarises the relationship between seizure dissimilarities and temporal distances across all patients. In almost all patients, there was a positive Spearman's corre-

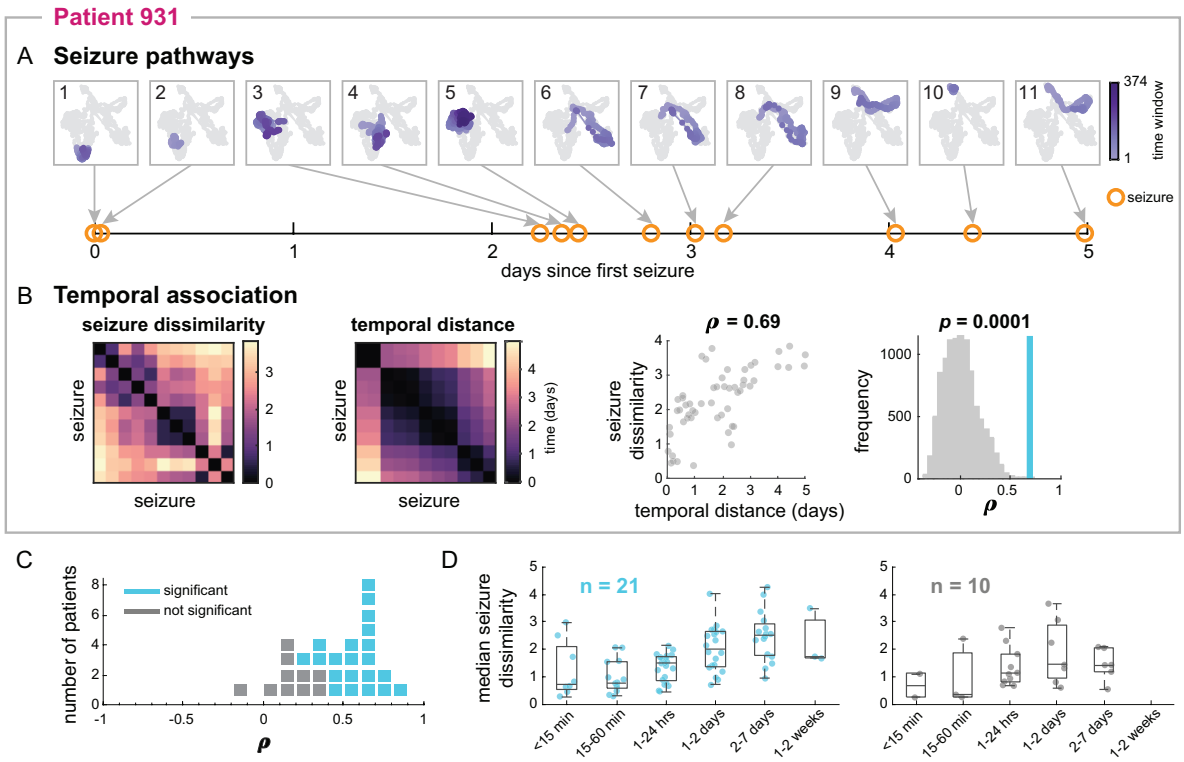

Figure 4.1: **More similar seizures tend to occur closer together in time in most patients.** A) Multidimensional scaling (MDS) projections of all of patient 931's seizure pathways, numbered from first to last seizure. The pathway of each seizure is shown in purple, with earlier time windows in lighter purple. In each plot, the pathways of the remaining seizures are shown in grey for comparison. Below the pathways, the time of each seizure (orange circles) relative to the first seizure is shown. B) From left to right: patient 931's seizure dissimilarity matrix, temporal distance matrix, and comparison of seizure dissimilarities and temporal distances. The seizure dissimilarity matrix captures differences in the patient's seizure pathways, while the temporal distance matrix quantifies the amount of time between each pair of seizures, in days. Plotting the seizure dissimilarity vs. the corresponding temporal distance of each pair of seizures (scatter plot, second from right) reveals a positive Spearman's correlation ρ between the two features. The significance (p-value) of this correlation was determined using a permutation test (far right panel); the observed correlation (blue) is shown against the distribution of correlations from permuted matrices (grey). C) Dot plot showing the range of correlations between seizure dissimilarities and temporal distances across all patients. Each marker represents a patient (blue = significant correlation, grey = not significant after false discovery rate correction). D) Median seizure dissimilarities of pairs of seizures occurring within different time intervals (i.e., temporal distances) for patients with (left, blue) and without (right, grey) a significant correlation between seizure dissimilarities and temporal distances. Each point corresponds to the median dissimilarity of pairs of seizures occurring within the given time interval in a single patient. Some time intervals have fewer observations since some temporal distances were not observed in all patients. The boxplots indicate the minimum, lower quartile, median, upper quartile, and maximum of the distribution of median seizure dissimilarities, across the subset of patients, for that time interval.

lation between seizure dissimilarities and temporal distances (range: -0.10 – 0.83, mean: 0.45). This association was significant in 21 patients (67.7%) after false discovery rate correction. In these patients, we also observed that the average level of dissimilarity tended to increase with the time between the two seizures (Fig. 4.1D). Interestingly, there was no association between whether antiepileptic medication was reduced and whether the correlation between seizure dissimilarities and temporal distances was significant (χ^2 test, $p = 0.9$) (Supplementary Section 4.5.1). Therefore, although medication levels may affect seizure occurrence and dynamics (Engel and Crandall, 1983; Marciani and Gotman, 1986; Napolitano and Orriols, 2013; Spencer et al., 1981), medication changes alone could not explain the observed shifts in seizure pathways, suggesting that other factors also play a role in shaping seizure features.

4.2.2 Changes in seizure pathways on different timescales can be modelled by a combination of slow, circadian, and noisy factors

The observed temporal associations of seizure dissimilarities reflected gradual changes in seizure network evolutions across the length of each recording. In other words, we observed relatively slow shifts in seizure pathways over the course of multiple days. We also hypothesised that seizure pathways may change on shorter timescales due to, for example, circadian rhythms. We first used a similar approach to Section 4.2.1 to determine if seizures that occurred at similar times of day had more similar pathways (Supplementary Section 4.5.2). However, contrary to our expectations, in most patients there was little association between seizure dissimilarities and "circadian differences" (i.e., the differences in the times-of-day of seizure occurrences). Across patients, the mean correlation between these measures was close to zero (mean: 0.05, range: -0.40 - 0.51), and the association was only significant in two patients after false discovery rate (FDR) correction for multiple comparisons.

We next hypothesised that circadian influences on seizure pathways may be obscured by the strong gradual shifts observed in Section 4.2.1. We therefore first computed a measure that characterised multiscale temporal relationships between seizure pathways. One of the drawbacks of the approach in Section 4.2.1 is that the overall correlation between seizure

dissimilarities and temporal distances only measured associations between seizures over the timescale of the patient's entire recording (Legendre and Gauthier, 2014). However, faster rhythms in seizure variability would also create timescale-dependent relationships between seizures; in particular, circadian changes would create a positive correlation between seizure dissimilarities and temporal distances on shorter timescales, but this association would be destroyed over longer timescales. Therefore, to explore the possibility of different timescales of changes in seizure pathways, we scanned the correlation between seizure dissimilarities and temporal distances on different timescales T ranging from 6 h to the longest amount of time between a seizure pair (Fig. 4.2A). For example, for $T = 3$ days, we computed the correlation between seizure dissimilarities and temporal distances for all pairs of seizures that occurred within three days of each other. We refer to this set of correlations as a "temporal correlation pattern" of seizure pathways. This approach is also known as a correlogram in ecological research (Legendre and Gauthier, 2014) and can be used to uncover temporal autocorrelations over different timescales.

Fig. 4.2A shows the temporal correlation pattern of patient 931's seizures. As we determined earlier, there was a positive correlation between seizure dissimilarities and temporal distances when all seizures were included in the computation ($T = 5$ days) as a result of the observed gradual changes in seizure pathways. At shorter timescales, however, the temporal relationship fluctuated; for example, the correlation was relatively low at $T = 1$ and 2.5 days, and higher at $T = 0.75$ and 2.5 days. These fluctuations were signs of additional, timescale-dependent changes in seizure pathways.

To investigate how these temporal correlation patterns arose, we modelled different patterns of seizure variability and the corresponding temporal correlation patterns (see Section 4.4.2 and Supplementary Sections 4.5.3 through 4.5.6 for modelling details). For each patient, we then determined which pattern(s) of changes were most likely to reproduce the observed dynamics. In particular, we classified patients as having

1. linear changes in seizure pathways (Fig. 4.2B, left column), which corresponded to the slower, gradual shifts in seizure evolutions,

2. circadian changes (Fig. 4.2B, middle column), which represented dynamics modu-

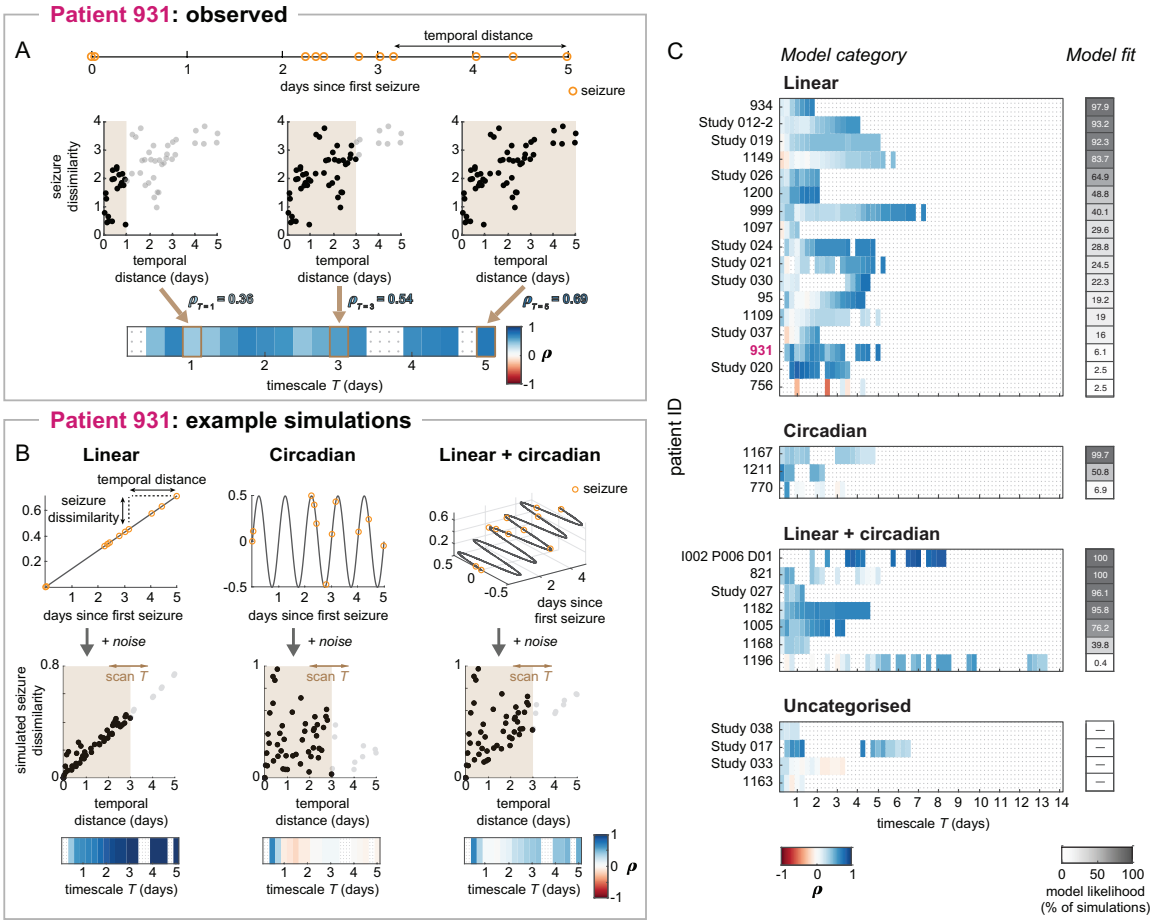

Figure 4.2: **Temporal patterns of changes in seizure pathways.** A) Calculation of patient 931's temporal correlation pattern. Temporal distances between seizures, which were derived from the patient's seizure times (top), were compared to seizure dissimilarities at different timescales (middle row). Example timescales $T = 1, 3, and 5$ days are shown (scatter plots, middle row). In each scatter plot, brown shading indicates the timescale, black points correspond to seizure pairs used to compute the correlation for that timescale, and grey points were pairs excluded from the correlation computation. At $T = 5$ days, all seizure pairs are included, producing the same temporal correlation as in Fig. 4.1B. Scanning the timescale produces a set of correlations, or "temporal correlation pattern," shown in the heatmap (bottom). Grey dots in the heatmap indicate insufficient information at that timescale, and these timescales are excluded from downstream analysis. B) Seizure dissimilarities were modelled based on linear (left column), circadian (middle column) or a combination of linear + circadian (right column) changes in seizure pathways. The simulated changes in seizure pathways are shown as different functions of time (top row), with patient 931's seizures marked in orange. Temporal distances and simulated seizure dissimilarities were compared across different timescales (scatter plots, middle row), yielding a temporal correlation pattern for each model (heatmaps, bottom row). Example models and simulation results are shown here; the full set of tested models is provided in Supplementary Section 4.5.3. C) Observed temporal correlation patterns of seizure pathways in each patient, categorised by the model that best reproduces these dynamics. The goodness of model fit was measured using "model likelihood" (grey heatmap). The full details of each patient's model are provided in Supplementary Section 4.5.5

lated by daily rhythms, or

3. some combination of the both linear and circadian changes (example combination shown in Fig. 4.2B, right column).

Fig. 4.2B demonstrates how seizure dissimilarities were simulated using patient 931's seizure times and example models from each of the above categories. In each model (Fig. 4.2B, top row), the x-axis value of each seizure gives the seizure's time, relative to the first seizure. These values are the same across all three models because they are the empirically observed seizure times of patient 931. Thus, the x-axis distance between a pair of seizures measures the amount of time, or temporal distance, between them. Based on the seizure times, each model then predicted how seizure pathways would change from seizure to seizure; specifically, the distance between two seizures along the other dimension(s) corresponds to the simulated dissimilarity of each pair of seizures. Each model additionally included noisy dynamics that allowed for further, random fluctuations in seizure pathways and thus seizure dissimilarities (Supplementary Section 4.5.4). By incorporating both circadian and linear changes in this model, we could identify more subtle circadian changes that would have been obscured by strong slow (linear) changes in seizure pathways and *vice versa*. Incorporating noise in the model and analysing the results of different noise realisations also allowed us to reproduce seizure variability patterns that were not strictly due to circadian/linear changes and determine how often such patterns could arise due to random effects. We note that simulating the model's effects on seizures directly, as in the top row of Fig. 4.2B (rather than fitting the model directly to seizure dissimilarities and temporal distances) allows us to more easily account for interdependencies present in the pairwise measures. For example, the effect of noise on one seizure's pathway will impact that seizure's similarity to all other seizures, as in our model.

From the temporal distances and simulated seizure dissimilarities (Fig. 4.2B, middle row), we then computed the corresponding temporal correlation patterns (Fig. 4.2B, bottom row) using the same process shown in Fig. 4.2A. A linear change in seizure pathways produced a positive temporal relationship that was stronger at longer timescales. Meanwhile, a circadian model only produced strong, positive temporal correlations at timescales

shorter than one day. Finally, a combination of the linear and circadian factors created both the short-term temporal relationships and a positive temporal correlation at the longer timescales. Note that there were also some additional fluctuations in the temporal correlation patterns due to noisy changes in dynamics; these effects differed depending on the outcome of the noisy simulation. To fully explore these noisy effects, we therefore additionally varied the level of noise added to the models. The tested combinations of noisy, linear, and circadian contributions are provided in Supplementary Section 4.5.3. For each combination of these factors, we simulated temporal correlation patterns 1000 times using different noise realisations to produce a series of possible temporal correlation patterns for each model.

Fig. 4.2C shows the type of model (linear, circadian, or linear + circadian) most likely to underlie the observed temporal correlation pattern of each patient. As a measure of model fit, we defined "model likelihood" as the percentage of model simulations that reproduced the patient's observed temporal correlation pattern. Model likelihood ranges from 0-100%, and higher values reveal that the modelled changes in dynamics consistently produced temporal correlation patterns that were similar to the patient's observed temporal correlation pattern. We additionally required the selected model to 1) outperform noisy simulations alone, 2) clearly distinguish between the linear and circadian models, and 3) in the case of the linear + circadian model, clearly outperform one of the simpler models. Using these criteria, 17 patients' temporal correlation patterns were best explained by the linear model, three by the circadian model, and seven by the linear + circadian model. Thus, most patients (77.4%) required a linear component to explain the observed changes in seizure pathways, while 32.3% of the patients were well-matched by a model incorporating circadian changes in pathways. Notably, model likelihood tended to be higher for patients with higher number of seizures, reflecting greater model certainty in cases with larger sample sizes (Supplementary Section 4.5.6). These different classifications of seizure variability were not associated with surgical outcomes (Schroeder et al., 2020) or whether the patient's antiepileptic medication was reduced during presurgical monitoring (Supplementary Section 4.5.7).

Notably, two of the three patients that were best fit by the circadian model, patients 1167

and 1211, were the two patients with significant associations between seizure dissimilarities and circadian distances (Supplementary Section 4.5.2)). The third patient, patient 770, was not identified by the circadian distance analysis, but this patient also had a weaker fit to the circadian model (model likelihood = 6.9), revealing that their temporal correlation pattern also relied on specific patterns of noisy fluctuations in seizure pathways. Thus, our two circadian analysis approaches agreed in the patients with the strongest circadian changes in seizure pathways, and this correspondence supports the ability of our model to uncover circadian patterns in seizure variability.

Four patients' temporal correlation patterns could not be assigned to a model, either because the linear model and circadian model performed similarly (patient Study 038) or the best model did not outperform noise alone (patients Study 017, Study 033, and 1163). Additionally, in some patients (e.g., patients Study 020, 756, and 1196) only a small percentage of the simulations matched the observed temporal correlation patterns, indicating that reproducing the observed dynamics required specific patterns of noise. In these cases, other models may therefore provide a better explanation for the patient's changes in seizure pathways. In particular, many of these patients had strong positive correlations at a timescales longer than one day, but less than the length of the recording, suggesting multi-day fluctuations in seizure pathways.

4.3 Discussion

In this chapter, we have explored how seizure pathways, defined as seizure network evolutions, change over time within individual patients with focal epilepsy. Using our quantitative measure of seizure dissimilarity, we have shown that more similar seizures tend to occur closer together in time. Our modelling results further indicate that in most patients, a combination of circadian and/or slower changes in seizure pathways may underlie the observed variability, suggesting that factors operating on different timescales modulate within-patient seizure evolutions.

One limitation of our work is that we analysed short-term recordings, ranging in length from a couple of days to a couple of weeks. As such, we were unable to observed cycles in seizure pathways over longer timescales (Baud et al., 2021). Analysing prolonged recordings of patients with focal epilepsy (Cook et al., 2013; Jarosiewicz and Morrell, 2021) may provide further insight into how pathological brain dynamics change over time and influence seizure features. In particular, recent studies using such data have shown that the rates of epileptiform discharges and seizures fluctuate according to both circadian and patient-specific multidien (approximately weekly to monthly) cycles (Baud et al., 2018; Karoly et al., 2018a, 2021; Leguia et al., 2021). An intriguing possibility is that the same factors that rhythmically modulate seizure likelihood may also influence seizure evolution. Consistent with this hypothesis, we found that the majority of observed temporal patterns of seizure variability were well-explained by models incorporating circadian and/or linear changes in seizure pathways. In particular, the linear component of the model may reflect gradual changes in pathways on slower timescales, ranging from weeks to months. These simple models provided an initial hypothesis for the observed patterns of changes in seizure evolutions. Some patients' seizure patterns may be better explained by more complex models that capture different dynamics, such as multistability or multidien cycles. Ultimately, it is likely that various factors, with differential effects on seizure evolution, interact to produce the observed repertoire of seizure pathways. Analysing within-patient seizure variability in long-term recordings could provide additional insight into such patterns of changes in seizure pathways.

Many of the patients in our study underwent antiepileptic medication reduction as part of pre-surgical monitoring, making it difficult to disentangle the effects of changing drug levels from other potential slow-varying modulators of seizure pathways. Changes in antiepileptic medication can impact neural excitability (Badawy et al., 2010, 2013; Meisel et al., 2016, 2015), and medication tapering increases seizure likelihood in most patients (Bardy, 1992; Marciani and Gotman, 1986); however, it is controversial whether it also affects seizure patterns (Bardy, 1992; Engel and Crandall, 1983; Marciani and Gotman, 1986; Spencer et al., 1981). In some cases, it appears that medication tapering reveals latent seizure pathways that are suppressed by medication (Spencer et al., 1981) or allows existing pathways to further progress (e.g., the secondary generalisation of typically focal seizures) (Marciani and Gotman, 1986). It is possible that the impact of medication reduction on seizure dynamics is drug-, patient-, and dose-dependent, and may ultimately depend on how well the medication controls neuronal excitability (Napolitano and Orriols, 2013). However, medication changes alone cannot account for the observed seizure variability in our cohort, as we observed temporal associations of seizure pathways in patients that did not undergo medication reduction. In future work, associating medication levels with differences in seizure pathways could help untangle the different factors shaping seizure dynamics.

Another confounding factor in our data is that the surgical implantation itself could artificially alter seizure dynamics over the length of the recording. Using chronic recordings of epileptic canines, Ung et al. (2016) found variability in seizure onset and interictal burst dynamics, with the most stable dynamics emerging approximately a few weeks after electrode implantation. In agreement with their work, we found that earlier seizure evolutions often recurred later in the recording, making it unlikely that gradual changes in the recording quality or an acute reaction to the surgery underlay the observed variability. Instead, Ung et al. (2016) hypothesised that seizure variability results from transient, atypical dynamics as the brain recovers from surgery, with later dynamics representing a truer epileptic network. Other stressors, such as medication withdrawal, could similarly elicit abnormal dynamics. Nevertheless, a large number of our patients had good surgical outcomes, suggesting that their recorded seizures accurately represented their epileptic

networks. Additionally, clinicians often note that patients have typical seizures during iEEG recordings, as compared to preimplantation reports, despite the effects of surgery and medication withdrawal (Marciani and Gotman, 1986). As such, the observed seizure dynamics in our cohort may be part of their usual repertoires of seizure evolutions, even if some dynamics are only elicited by strong stressors. Further analysis in chronic human recordings, such as the NeuroVista dataset (Cook et al., 2016, 2013; Karoly et al., 2018b), is needed to determine whether and how seizure pathways vary in a more naturalistic setting.

More generally, our work adds to the growing literature on within-subject temporal variability in brain dynamics and other physiological states (Poldrack et al., 2015) in both health and disease. In particular, there is increasing interest in developing improved, time-sensitive treatments that adapt to the patient's changing state. Rather than delivering a continuous or regular therapy, such treatments would be modified and/or timed to improve their efficacy while also reducing treatment side effects. Treatment parameters may be tuned to biological rhythms (Baud et al., 2018; Baud and Rao, 2018; Karoly et al., 2018a; Navis and Harden, 2016; Ramgopal et al., 2013; Scheiermann et al., 2018) or respond directly to fluctuating conditions within each patient (Fava et al., 2003; Sanchez-Morillo et al., 2017; Vettoretti and Facchinetti, 2019). To investigate temporal fluctuations within each patient and determine how treatments interact with these changes, researchers may draw inspiration from spatiotemporal analyses in other fields, such as ecology (Legendre and Gauthier, 2014), genetics (Wu et al., 2016), and engineering (Thomas et al., 2019; Wang et al., 2015), as well as develop new techniques that address specific data-analytical challenges.

In summary, we have shown that seizure network evolutions do not randomly fluctuate within individual patients with focal epilepsy. Instead, patterns of temporal changes in seizure evolutions suggest that a combination of circadian and slow-varying factors shape these seizure pathways, perhaps by modulating the background brain state. Determining the factors that influence neural dynamics and seizure pathways could allow researchers and clinicians to develop new, seizure-specific treatments for altering seizure pathways and preventing seizure onset. Ultimately, understanding temporal patterns of changes in

seizure pathways could allow time-varying treatments to target specific seizure dynamics within each patient, ultimately improving patient outcomes and reducing treatment side-effects.

4.4 Methods

This chapter is a continuation of the analysis from Chapter 3; it analyses seizure network evolutions from the same 31 EMU patients using the methods outlined in that chapter. Additional methods used to analyse temporal patterns in seizure network evolutions are outlined here.

4.4.1 Comparison to temporal distances

For each patient, we computed a "temporal distance matrix" containing the amount of time elapsed (measured in days) between the onset times of each pair of seizures. Spearman's correlation was computed between the upper triangular elements of the seizure dissimilarity matrix and the temporal distance matrix of each patient. Since the distances in each matrix were not independent observations, the Mantel test (Mantel, 1967) was used to determine the significance of each correlation. Briefly, the rows and columns of one matrix were randomly permuted 10,000 times. The correlation between the two sets of upper triangular elements was re-computed after each permutation, resulting in a distribution of correlation values that described the expected correlation if there were no relationship between seizure dissimilarities and temporal distances. The p-value of the association was then defined as the proportion of permuted correlation that were greater than or equal to the observed correlation. To correct for multiple comparisons, the Benjamini-Hochberg FDR correction (Benjamini and Hochberg, 1995) was applied to the set of p-values computed across all patients (31 total tests). The correlation was considered significant if the corresponding adjusted p-value was less than 0.05.

4.4.2 Computing and modelling temporal correlation patterns

To quantify how seizure dynamics change over different timescales in each patient, Spearman's correlation between seizure dissimilarities and temporal distances was computed only for seizure pairs with temporal distances less than or equal to timescale T. T was scanned from 0.25 days up to the patient's largest temporal distance in steps of 0.25 days. A timescale was excluded from the analysis if less than seven pairs of seizures occurred

within the given timescale or if no new seizure pairs were added when the timescale was increased. The resulting set of correlations across various timescales were referred to as "temporal correlation patterns."

Modelling seizure dissimilarities and temporal correlation patterns: To determine the underlying processes that could produce the observed temporal correlation patterns, changes in seizure dynamics were modelled using the functions

- $f_l(t) = 1/7t$ (a line with a slope of one per week)
- $f_c(t) = sin2\pi t$ (a sine wave with a period of one day)
- $f_n(t) \sim N(0, 1)$ (Gaussian noise with a mean of zero and standard deviation of 1)

where t is time in days.

For each function, a simulated distance matrix D was then defined for the patients' seizures, with

$$D(i, j) = |f(t_i) - f(t_j)|$$

where t_i is the time of seizure i, t_j is the time of seizure j, and $f(t)$ is the corresponding function. The simulated dissimilarity of the two seizures was then defined as

$$Diss(i, j) = \sqrt{[lD_l(i, j)]^2 + [cD_c(i, j)]^2 + [nD_n(i, j)]^2}$$

where l, c, and n are scalar parameters controlling the relative contributions of the linear, circadian, and noise functions, respectively.

The relative contributions of the linear, circadian, and noise functions were scanned by varying the levels of l, c, and n. For each set of parameters, seizure dissimilarities were simulated 1000 times using different noise realisations (and correspondingly changing the noise distance matrix, D_n), and the resulting temporal correlation patterns were computed for each set of simulated dissimilarities. Note that because temporal correlation patterns only depend on the order of the dissimilarities, only the relative magnitudes of

the parameters l, c, and n affected the modelling results. A model was termed a "linear model" if $c = 0$, a "circadian model" if $l = 0$, and a "linear + circadian model" if $l > 0$ and $c > 0$.

To determine if a patient's seizure dynamics could be categorised as linear, circadian, or linear + circadian, the simulated temporal correlation patterns were compared to the patient's observed temporal correlation pattern by computing the mean squared error (MSE) of each simulated pattern. Simulated temporal correlation patterns with MSE ≤ 0.02185 were defined as "good matches" to the observed dynamics. This threshold was chosen because it was the 5th percentile of the set of all MSEs, across all patients, and based on visual inspection of simulated temporal correlation patterns with different MSEs. The likelihood L of a given parameter set was then defined as the percentage of "good matches" produced by the 1000 noisy simulations of seizure dissimilarities at those parameter values. For each class of model (linear, circadian, or linear + circadian), the model's likelihood (L_l, L_c, or L_{l+c}, respectively) was the highest likelihood among the model type's parameter sets, and the "best model" was the model with the highest likelihood. L_n was also defined as the highest likelihood of the parameter sets without any linear or circadian contributions ($l = 0$, $c = 0$, $n > 0$).

This best model with likelihood L_{max} was then used to categorise the patient's dynamics if it outperformed all competing models. Specifically, we required that

- The best model clearly outperform noise alone ($L_{max} \geq 2L_n$); otherwise, the patient's dynamics were classified as other/indeterminate.

- The performance of the linear model and circadian model were clearly distinguishable ($L_l \geq 2L_c$ if the linear model was best; $L_c \geq 2L_l$ if the circadian model was best); otherwise, the patient's dynamics were classified as other/indeterminate.

- If the best model was linear + circadian, it clearly outperform the two simpler models ($L_{l+c} \geq 2L_l$ and $L_{l+c} \geq 2L_c$); otherwise, the patient's dynamics were classified as the simpler model (if one simpler model performed comparably by this criterion) or as other/indeterminate (if both simpler models performed comparably).

See Supplementary Sections 4.5.3 through 4.5.6 for additional modelling details and the selected models for each patient.

4.5 Supplementary

4.5.1 No relationship between temporal correlation and AED reduction

During presurgical recordings, the antiepileptic drug (AED) dosages of patients are often gradually reduced to provoke more seizures and thus facilitate localization of the epileptogenic zone (Bardy, 1992; Engel and Crandall, 1983; Marciani and Gotman, 1986; Spencer et al., 1981). Because such medication alters neural excitability (Badawy et al., 2010; Meisel et al., 2016, 2015), it is possible that it also affects seizure dynamics. Thus, gradual changes in AED dosages could potentially produce gradual changes in seizure evolution, leading to the observed temporal changes in seizure pathways in our work. We therefore investigated if differences in the correlation between seizure dissimilarities and temporal distances across patients could be attributed to whether antiepileptic medications were reduced in each patient.

Information about medication dosages was available from the clinical reports of 19 patients. This information was extracted and used to label each patient as "AED reduction performed" or "no AED reduction performed." Medication changes due to stat doses of medication (i.e., medication given in addition to the planned dose in order to control seizures) were not considered in this assessment; rather, we sought to identify whether gradual changes in medication were intentionally made, as such changes could underlie the observed gradual temporal changes in seizure pathways. We then also labelled whether each of the 19 patients had a significant or not significant correlation between seizure dissimilarities and temporal distances.

Table 4.1 shows the cross-tabulation table of these variables. If AED reduction reliably altered seizure evolution, we would expect significant correlations between seizure dissimilarities and temporal distances in approximately all patients who underwent medication reduction. However, of the 11 patients that underwent AED reduction, four did not have a significant association between these distances. Additionally, if the temporal changes in seizure pathways were solely due to medication reduction, we would not expect a significant relationship between seizure dissimilarities and temporal distances in patients who did not undergo AED reduction. Instead, we observed a significant correlation in five of

		AED reduction performed		
		no	yes	total
Significant temporal correlation	no	3	4	7
	yes	5	7	12
	total	8	11	19

Table 4.1: **Cross-tabulation table of AED reduction and significant temporal correlations.** The frequency of AED reduction (columns, "AED reduction performed") and the frequency of a significant correlation between temporal distances and seizure dissimilarities (rows, "Significant temporal correlation") in our cohort of patients.

the eight patients who did not undergo AED reduction. A χ^2 test found no association between whether AED reduction was performed and whether the correlation between seizure dissimilarities and temporal distances was significant ($\chi^2 = 0.0026, p = 0.96$). As such, the temporal relationship between similar seizure pathways cannot solely be attributed to AED reduction in our cohort.

4.5.2 No relationship between seizure dissimilarities and circadian timing alone in most patients

In many patients with focal epilepsy, seizure likelihood (Baud et al., 2018; Karoly et al., 2017; Leguia et al., 2021; Maturana et al., 2020; Stirling et al., 2021) and seizure dynamics (Bazil, 2018; Bazil and Walczak, 1997; Sinha et al., 2006) vary according to circadian timing and cycles. We therefore investigated if within-patient variability in seizure pathways could be explained by differences in the circadian timing of the seizures.

For each patient, we computed a circadian distance matrix, which contained the differences (measured in hours) in the time-of-day of the occurrence of each pair of seizures. This measure is a circular statistic that can range from 0 to 12 hours; for example, the circadian distance between two seizures that occur at 1 am and 11 pm is two hours, not 22 hrs. As for the temporal distance matrix (section 4.4.1), Spearman's correlation was computed between the upper triangular elements of a patient's seizure dissimilarity and circadian distance matrices to determine if seizures with more similar circadian timings had more similar pathways, and the Mantel test (Mantel, 1967) was used to determine the significance of the correlation. To correct for multiple comparisons, the Benjamini-

Hochberg false discovery rate (FDR) correction (Benjamini and Hochberg, 1995), with $\alpha = 0.05$, was applied to the set of p-values across all patients.

Fig. 4.3 shows this comparison in three patients, as well as the correlations between seizure dissimilarities and circadian distances across patients. In patient 931 (Fig. 4.3A-D), there was no relationship between seizure dissimilarities and circadian distances; indeed, the correlation between these measures was close to zero. Of our cohort, patient 1211 (Fig. 4.3E-H) had the highest correlation between seizure dissimilarities and circadian distances; however, most seizures occurred within the same 24 hour period in this patient; thus, the observed changes could also be due to slower shifts in seizure pathways. We therefore also show this relationship in patient 1167 (Fig. 4.3I-L), whose seizures spanned multiple days. This patient had a significant association between seizure dissimilarities and circadian distances, but the relatively low correlation and the visual differences between the two matrices also suggest that much of the variability in seizure pathways was unexplained by the circadian timings of the seizures.

Across patients, the distribution of correlations between seizure dissimilarities and circadian distances was centred close to zero (mean: 0.05, range: -0.40 – 0.51), with weak associations in most patients (Fig. 4.3M). Further, the correlation was significant in only two patients. Contrary to our expectations, these results indicate that the circadian timing of seizures alone does not explain the observed variability in seizure network dynamics in our cohort. It is possible that only some seizure features, such as secondary generalisation (Bazil, 2018; Bazil and Walczak, 1997; Sinha et al., 2006), are influenced by circadian rhythms. Additionally, seizure features may be more influenced by wakefulness and sleep stage than by time-of-day; for example, seizures that have relatively similar circadian timing (e.g, 7 am and 8 am) could be highly dissimilar if they span a transition from wake to sleep. Future studies could detect sleep and wake states from iEEG recordings (Conrad et al., 2020) to test this hypothesis.

The weak relationship between seizure dissimilarities and circadian distances could also be due to other, stronger factors that influence seizure pathways on different timescales, obscuring more subtle circadian effects. Indeed, many of our patients had strong shifts in seizure pathways across the timescale of their recordings (section 4.2.1). These slower

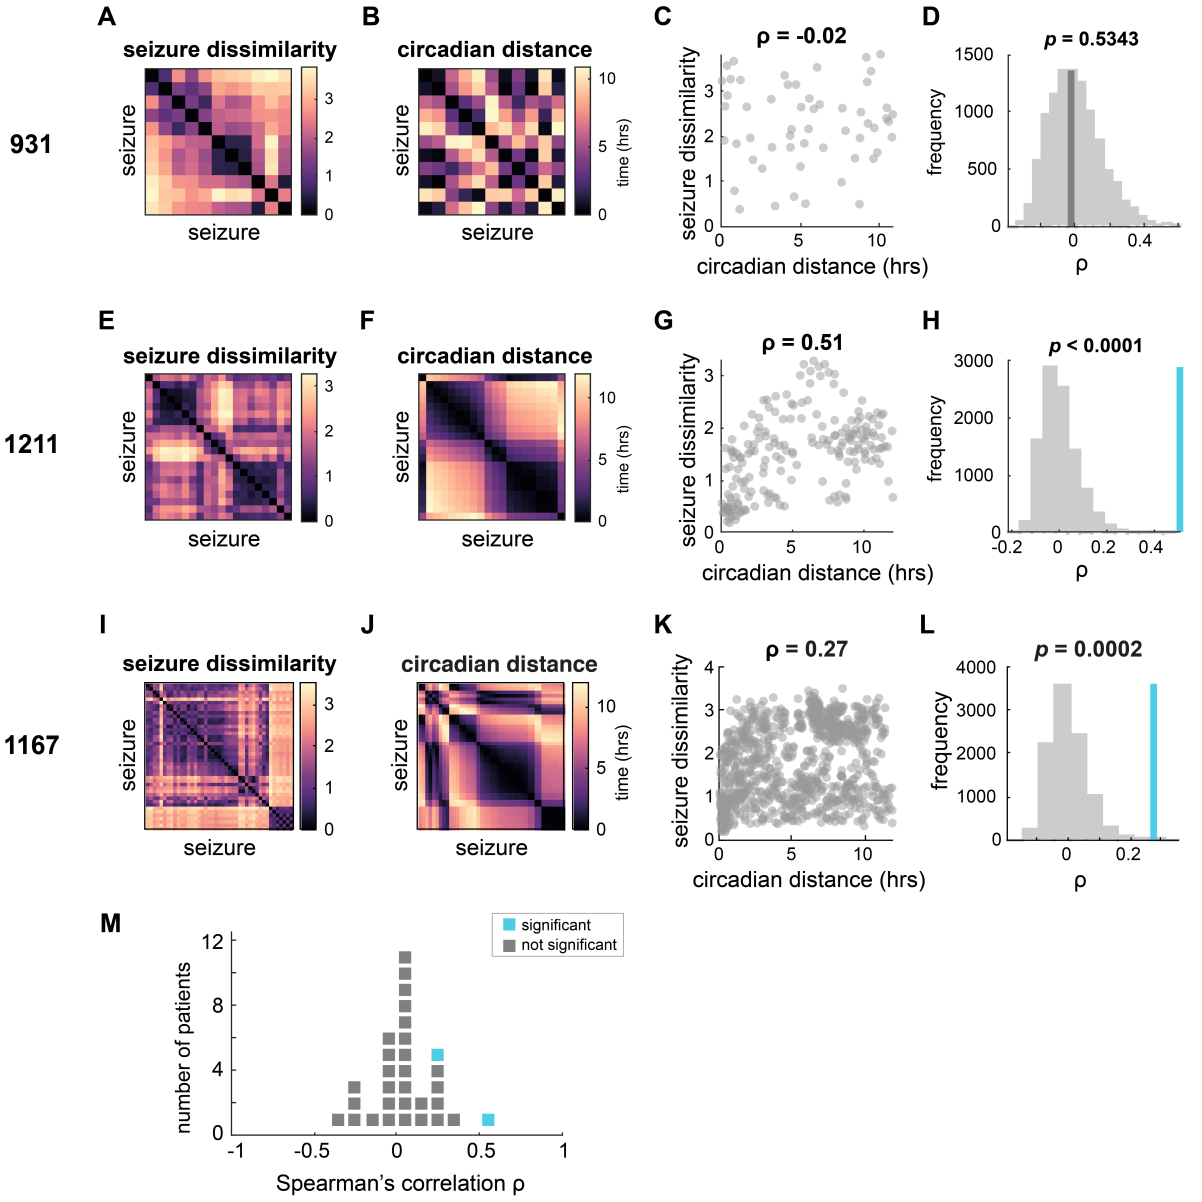

Figure 4.3: **Relationship between seizure dissimilarities and circadian distances.** For three example patients, patients 931 (A-D), 1211 (E-H), and 1167 (I-L): the patient's seizure dissimilarity matrix (A, E, I), circadian distance matrix (B, F, J), scatter plot of the seizure dissimilarities vs. circadian distances (C, G, K), and permutation tests results for determining the significance of the Speaman's correlation between seizure dissimilarities and circadian distances (D, H, L). M) Dot plot showing the range of Spearman's correlations between seizure dissimilarities and circadian distances across all patients. Each marker represents a patient (blue = significant correlation, grey = not significant after false discovery rate correction).

changes may need to be accounted for in order to observe circadian influences on seizure pathways. Thus, in section 4.2.2, we used a model that incorporated both slow (linear) and circadian factors to explain temporal patterns in seizure variability.

4.5.3 Model parameter scan

Table 4.2 shows the values used to scan the parameters l, c, and n in the model of seizure temporal correlation patterns. To reduce the model complexity, we do not change the phase of the circadian cycle; however, in future work, the circadian phase could also be included as an additional parameter. Alternatively, the parameter could be tuned based on the patient's known or likely sleep/wake cycle (data that was unavailable for our patients).

parameter	**description**	**scan range**	**step size**
l	scales linear contribution (i.e., slow/gradual changes)	0 to 1	1
c	scales circadian (sinusoidal) contribution	0 to 5	0.1
n	scales noise contribution	0.025 to 0.5	0.025

Table 4.2: **Parameter scan values for model of temporal correlation patterns.**

Note that temporal correlation patterns only depend on the relative magnitude of seizure dissimilarities, and not their absolute magnitude, since they are computed using Spearman's correlation. As such, only the relative values of the model parameters matter for model selection. For example, given the same noise realisation, the temporal correlation patterns of the parameter sets ($l = 0.5$, $c = 1$, $n = 0.05$) and ($l = 1$, $c = 2$, $n = 0.1$) would be equivalent. These properties allowed us to limit the parameter scan, removing redundant parameter sets, in the following ways:

- We limited the values of l to 0 (no linear contribution) or 1 (linear contribution), and only scanned the values of c and n, relative to these fixed linear contributions. Therefore, for all model that include a linear component, $l = 1$. The values of the other parameters then indicate the contributions of circadian and/or noisy changes in dynamics, relative to the contribution of the linear component.

- For parameter sets where $l = 0$, we likewise fixed the value of c to 1 and only scanned n to determine the relative contributions of the circadian and noise processes. Thus, for any patient whose pattern of dynamics was categorised as "circadian" (with no linear contribution), $l = 0$ and $c = 1$.

4.5.4 Effect of level of noise on modelled temporal correlation patterns

For the example models shown in Section 4.2.2, Fig. 4.2, we visualise how increasing the level of noise alters the temporal correlation pattern (Fig. 4.4). In general, increasing the level of noise attenuates the correlations produced by the linear and/or circadian dynamics. Additionally, specific realisations of noisy dynamics can produce random fluctuations in the correlation pattern (i.e., correlations that are much higher or lower than expected based on linear and/or circadian dynamics).

4.5.5 Modelling results of all patients and example computation of model likelihood

As a reminder, at each set of parameters, seizure dissimilarities and the corresponding temporal correlation patterns were simulated 1000 times. Each simulated temporal correlation pattern was then compared to the patient's observed temporal correlation pattern by computing the MSE between the simulated and observed patterns. The likelihood L of a given parameter set was defined as the percentage of "good matches" (simulations with MSE ≤ 0.02185) produced by the 1000 noisy simulations at those parameter values. A model was termed a "linear model" if $c = 0$, a "circadian model" if $l = 0$, and a "linear + circadian model" if $l > 0$ and $c > 0$. For each class of model (linear, circadian, or linear + circadian), the model's likelihood (L_l, L_c, or L_{l+c}, respectively) was the highest likelihood among the set of qualifying parameter sets, and the "best model" was the model with the highest likelihood, L_max. L_n was also defined as the highest likelihood of the parameter sets without any linear or circadian contributions ($l = 0$, $c = 0$, $n > 0$).

However, this "best model" was only the "selected" model for the patient if

1. The best model clearly outperformed noise alone ($L_{max} \geq 2L_n$); otherwise, the patient's dynamics were classified as other/indeterminate.

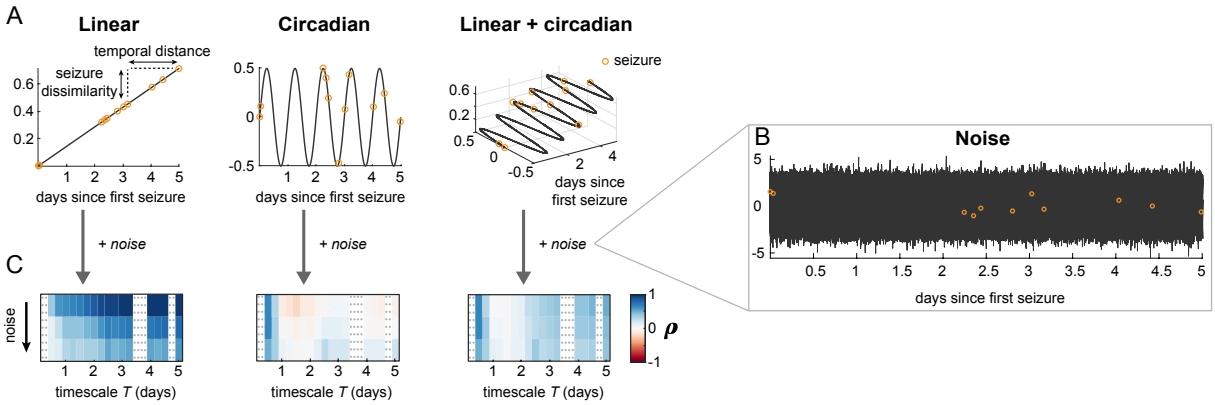

Figure 4.4: **Effect of noise level on example temporal correlation patterns.** Using example models of patient 931's seizure variability (reproduced from Fig. 4.2B), we demonstrate how increasing the level of noise impacts temporal correlation patterns. A) As in Fig. 4.2B, seizure dissimilarities were modelled based on linear, circadian, or varying combinations of linear + circadian changes in seizure pathways. In each model, distances between seizures reflected the observed temporal distances and the simulated seizure dissimilarities (see annotations on "Linear" model). B) In the same way, noisy changes in seizure dynamics were also incorporated into the models. The right panel shows the noisy dynamics (sampled from a standard Gaussian distribution) used in these example simulations; distances between seizures along the y-axis correspond to differences in seizure evolutions due to random fluctuations. C) The simulated seizure dissimilarities and temporal distances were compared across different timescales, yielding a temporal correlation pattern for each model (heatmaps, bottom row). For each model, the effects of three different levels of the additional, noisy dynamics are shown; specifically, from top to bottom of the heatmaps, the contribution of the noisy changes from panel B were multiplied by 0.1, 0.25, or 0.5 to vary their effect on the temporal correlation patterns. The top row of each heatmap is the temporal correlation pattern shown in Fig. 4.2B. Increasing the noisy contribution (second and third rows of the heatmaps) attenuates these correlations and/or produces additional fluctuations in the correlations.

2. The performance of the linear model and circadian model were clearly distinguishable ($L_l \geq 2L_c$ if the linear model was best; $L_c \geq 2L_l$ if the circadian model was best); otherwise, the patient's dynamics were classified as other/indeterminate.

3. If the best model was linear + circadian, it clearly outperformed the two simpler models ($L_{l+c} \geq 2L_l$ and $L_{l+c} \geq 2L_c$); otherwise, the patient's dynamics were classified as the simpler model (if one simpler model performed comparably by this criterion) or as other/indeterminate (if both simpler models performed comparably).

Fig. 4.5 shows the final modelling results for each patient after the above model selection criteria were applied. From the simulated temporal correlation patterns (Fig. 4.5B), it is apparent that our simple model can approximately reproduce the observed temporal correlation patterns (Fig. 4.5A). Fig. 4.5C provides the selected model parameters for each patient. Fig. 4.5D provides the likelihood of the selected model, and Fig. 4.5E shows the relative performances of the "best model," with L_{max}, compared to each of the model categories. The relative performance was computed as the likelihood of the best-performing model, L_{max}, divided by the likelihood of the given model category of model. Lower values indicate better performances compared to the best model; indeed, the relative performance of the best model is 1 because $L_{max}/L_{max} = 1$ (e.g., patient 934, linear model). These relative performances were used to select the final model parameters (Fig. 4.5C) and corresponding model category. For example, in some patients, the linear + circadian model performed best ($L_{max} = L_{l+c}$), but the linear model was selected because it performed comparably: in Study 012-2, for instance, $L_{l+c}/L_l = 1.058$, indicating that almost as many simulations from the linear model provided a good match to the observed temporal correlation patterns.

To illustrate how the model likelihood was computed, Fig. 4.6 shows the 1000 simulations arising from the selected parameter sets of three patients, 931, 1005, and 1211, which were modelled using the linear, linear + circadian, and circadian models, respectively. The simulations are ordered from lowest to highest MSE and are divided by whether they fall under the MSE threshold for a "good match" to their observed temporal correlation

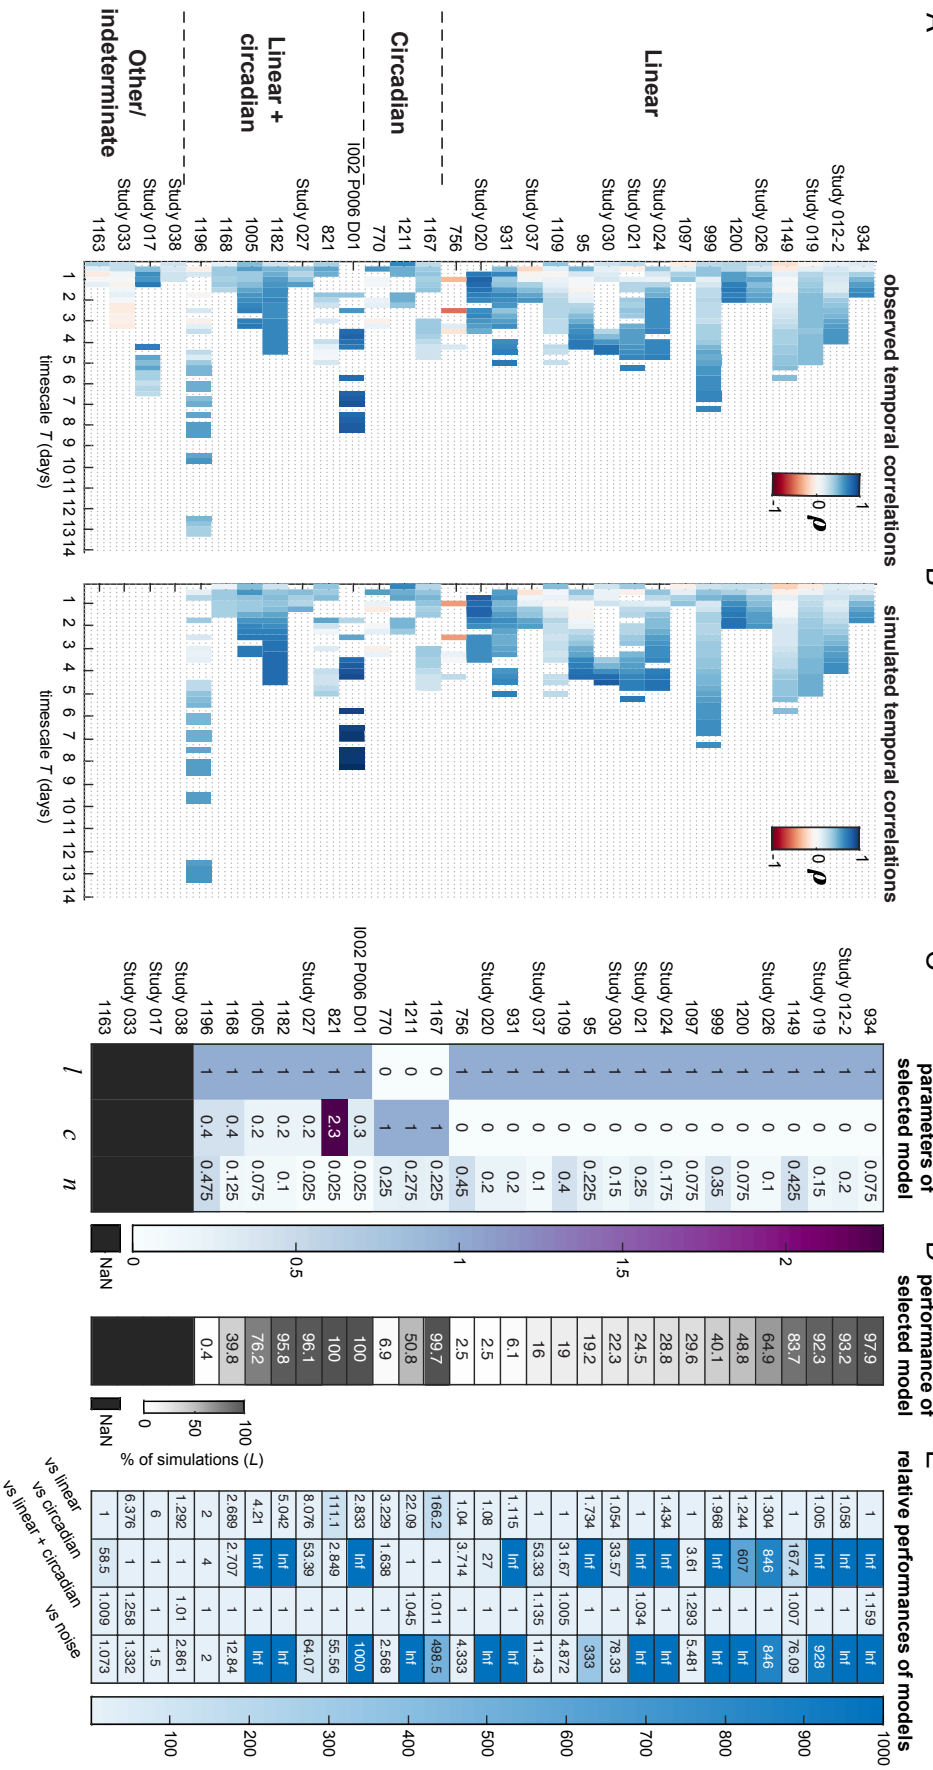
A
observed temporal correlations
B
simulated temporal correlations
timescale T (days)
Linear
Circadian
Linear + circadian
Other/ indeterminate
C
parameters of selected model
D
performance of selected model
E
relative performances of models
vs linear
vs circadian
vs linear + circadian
vs noise
% of simulations (L)
NaN

Figure 4.5: **Temporal correlation pattern modelling results for each patient** (previous page). A) Observed temporal correlation patterns of all patients, sorted by the selected model (linear, circadian, linear + circadian, or other/indeterminate). These results are also included in Fig. 4.2C. B) For each patient, the simulated temporal correlation pattern, arising from each patient's selected model parameters, that best matched the observed temporal correlation pattern. No simulated pattern is shown for the last four patients because none of the models (linear, circadian, or linear + circadian) met the model selection criteria. C) Parameter values of the selected model for each patient (l = linear contribution, c = circadian contribution, n = noise contribution). D) The performance of the selected model, as defined by the likelihood L of the selected model. A higher likelihood indicates that the model is more likely to produce the patient's observed temporal correlation pattern. E) The relative performances of the different model categories (linear, circadian, and linear + circadian, as well as the noise-only model) compared to the best-performing model. The relative performance is defined as the likelihood of the best-performing model, L_{max}, divided by the likelihood of the given type of model. Lower values indicate better performances compared to the best model. For each patient, the relative model performances were used to determine the model that clearly outperformed the other categories of models, while also providing the most parsimonious explanation of the observed dynamics.

pattern, which is shown in Fig. 4.5A. For each patient, the "good matches" are visually similar to the observed temporal correlation pattern of each patient. Note that the average MSE of the set of simulations would not necessarily be an appropriate evaluation for model selection because noisy fluctuations can dramatically alter the simulated temporal correlation pattern, especially if the number of analysed seizures was small. For example, for patient 931 (Fig. 4.6A), although some simulations provided a good match to the observed temporal correlation pattern, other simulations had a very high MSE. Our approach determined the "likelihood" of observing such good matches to the observed dynamics, without penalising the parameter set for also producing drastically different dynamics under different noise realisations.

4.5.6 Relationship between model likelihood, sample size, and the level of noise

Ideally, a measure of model fit reflects the amount of confidence in the model, and this confidence should increase with a larger sample size (here, the number of seizures). The model fit may also reflect how consistently the observed dynamics match the model dy-

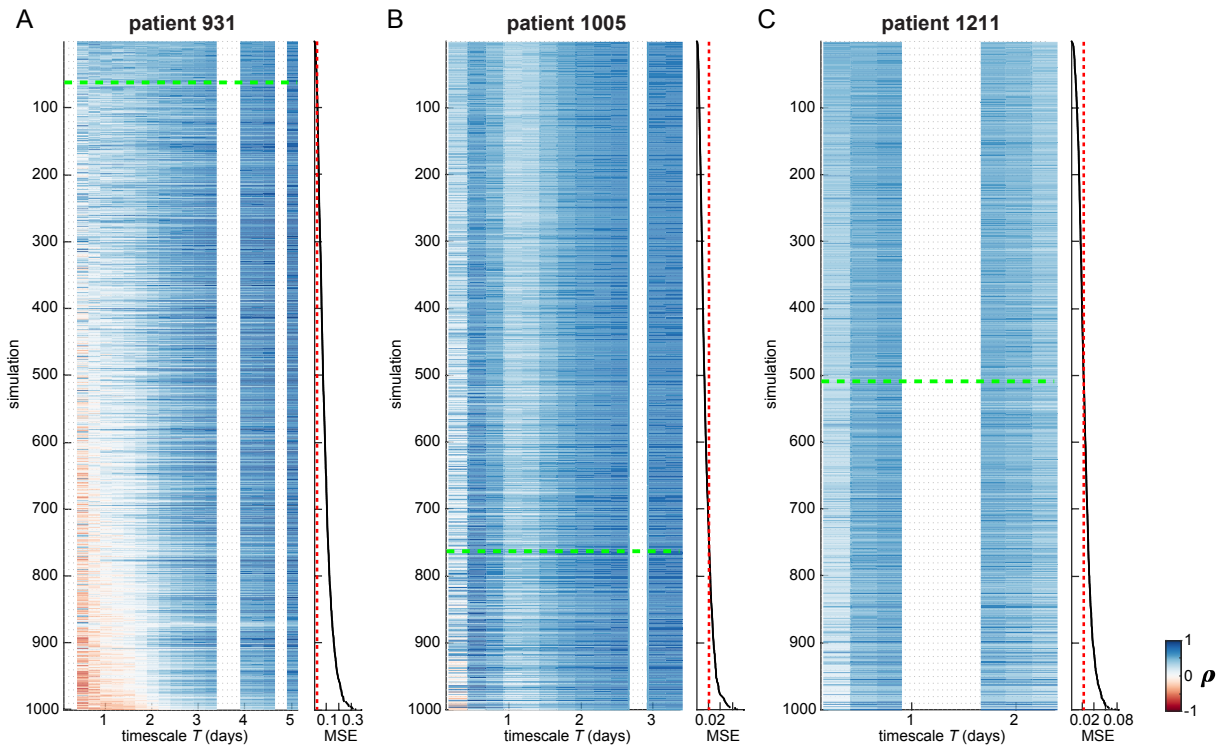

Figure 4.6: **All simulated temporal correlation patterns arising from the selected model parameters of three example patients:** A) patient 931, B) patient 1005, C) patient 1211. See Fig. 4.5C for the parameters used to generate the simulated temporal correlation patterns. From left to right, the example patients were categorised as linear, linear + circadian, and circadian. For each patient, the heatmap shows the temporal correlation patterns, with the MSE of each temporal correlation pattern (compared to the patient's observed temporal correlation pattern) to the right of the heatmap. Simulations are ordered from lowest MSE (top) to highest MSE (bottom). The red dotted line shows the MSE threshold for a "good match" to the observed temporal correlation pattern, and the dotted green line marks the boundary of the simulations that fall under this threshold.

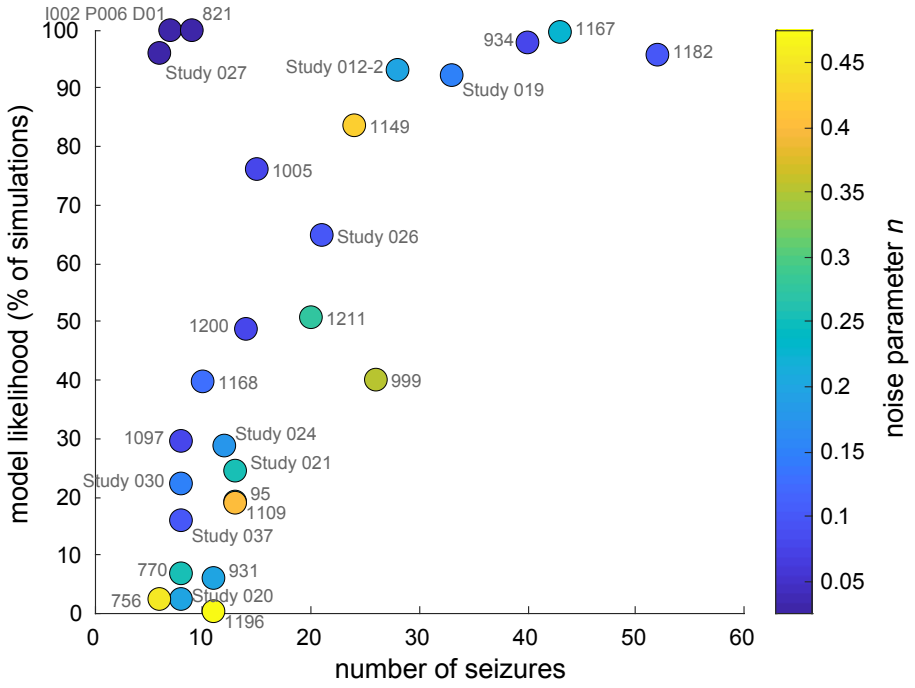

Figure 4.7: **Relationship between model likelihood, number of seizures, and the model noise parameter n.** Scatter plot of model likelihood (the percentage of model simulations, arising from the selected model parameters, that matched the patient's observed temporal correlation pattern) vs. the number of seizures analysed in each patient. Spearman's correlation between these two measures is 0.438. Points are labelled by the corresponding patient ID and coloured by the value of the noise parameter, n, in the selected model. See Fig. 4.5 for the exact values of model likelihood and n for each patient. Patients who were not assigned model parameters (patients Study 038, Study 017, Study 033, and 1163) are excluded from this figure.

namics. To evaluate our measure of model fit (model likelihood), we therefore compared the model likelihood of each patient to the number of seizures analysed in each patient (Fig. 4.7). We additionally visualised how the amount of noise in the selected model (the model parameter n) was related to model likelihood.

From Fig. 4.7, it is apparent that model likelihood tended to increase as the number of seizures analysed also increased. Spearman's correlation between these two measures was 0.438. Thus, our measure of model fit was sensitive to the number of seizures analysed, and we had more confidence in our model when the sample size was larger.

Notably, three patients with a relatively low number of seizures (patients Study 027, I002 P006 D01, and 821) had high model likelihood despite their small sample sizes. In these cases, the model likelihood was still high because their pattern of seizure variability was

very consistent with the linear and/or circadian changes predicted by the model; i.e., the noise contribution (controlled by the model parameter n) was low for all of these patients. The low n in turn meant that the selected model parameters created consistent simulated temporal correlation patterns that were all a good match for the observed dynamics. Therefore, even though the sample sizes for these patients was low, the model likelihood was high because the observed variability was highly consistent with the linear and/or circadian pattern of changes predicted by the model.

However, this relationship between model likelihood and the noise parameter n was only present for patients with low numbers of seizures. Patients with high numbers of seizures had high model likelihood regardless of the selected value of n. This situation arose because we compared observed and simulated correlations, rather than trying to predict specific fluctuations in seizure pathways; as such,

- When a patient had a low number of seizures, noisy fluctuations in seizure pathways could produce spuriously high or low correlations between seizure dissimilarities and temporal distances. Because these high correlations were unexplained by the linear or circadian trends in seizure dynamics, the model instead reproduced them by introducing higher levels of noisy dynamics. However, only a small number of these noisy simulations would reproduce the observed dynamics, resulting an overall low model likelihood (for e.g., see simulations of patient 931, with 11 seizures and $n = 0.2$, in Fig. 4.6A).

- In patients with a high number of seizures, spuriously high or low correlations become less likely to occur by chance. The same sample size effect would occur in the model, as well; due to the larger sample size, noisy fluctuations predictably attenuate the correlations created by other modelled changes. Thus, a high percentage of the noisy model simulations would match the observed dynamics, as long as those dynamics were also consistent with other modelled trends (i.e., the linear and/or circadian dynamics) (for e.g., see simulations of patient 1211, with 20 seizures and $n = 0.275$, in Fig. 4.6C).

Therefore, model likelihood does not necessarily reflect whether the observed dynamics

are consistent with the modelled linear/circadian changes in seizure pathways. Instead, such model uncertainty is directly incorporated into the model via the addition of noise. A higher contribution of noise reflects that additional changes in seizure pathways occurred due to other factors or random fluctuations.

Note that the relationship between model likelihood and the number of seizures only holds if the model is able to reproduce the observed pattern of seizure variability; models that cannot reproduce the observed temporal correlation patterns will always have a low or zero likelihood. Fig. 4.5E compares the performances of the different types of models, and it is apparent that model likelihood will be low, regardless of the sample size, if the model cannot reproduce the observed temporal correlation pattern. Therefore, model likelihood does not purely reflect the number of seizures, but also whether the model is a good fit for the observed dynamics. Thus, model likelihood is a useful measure for selecting models that reproduce the observed temporal correlation patterns in our cohort.

4.5.7 No relationship between model timescales and AED reduction

We additionally explored if patients who underwent AED reduction had seizure variability that was best described by a particular model (linear, circadian, or linear + circadian changes in dynamics). In particular, the linear model describes more gradual changes in seizure pathways over the course of the recording, which could be attributed to AED reduction.

As in Section 4.5.1, we analysed the 19 patients who had information about AED reduction in their clinical reports. Table 4.3 shows the number of patients with and without AED reduction who were assigned to each model category. A χ^2 test cannot be performed here due to small sample sizes and assignment to multiple categories (patients with linear + circadian variability could be considered members of both the linear and circadian groups). However, from the cross-tabulation table alone, it is apparent that similar proportions of patients with and without AED reduction were assigned to each model category. In particular, AED reduction alone cannot explain the linear pattern of seizure variability: six of the eight patients without AED reduction had variability that was categorised as either linear or linear + circadian. Thus, factors beyond AED reduction influence the

temporal patterns of changes in seizure pathways in this cohort.

	linear	circadian	linear + circadian	other/indeterminate
AED reduction	5	1	4	1
no AED reduction	4	2	2	0

Table 4.3: **Cross-tabulation table of model category and AED reduction.**

Chapter 5. Seizure Pathways and Seizure Durations Can Vary Independently Within Individual Patients with Focal Epilepsy

Contents

5.1 Introduction

Within individual patients with focal epilepsy, seizure spatiotemporal neural dynamics can vary in two main ways. First, the progression, or evolution, of pathological activity can differ from seizure to seizure. These progressions can be described with various features, such as functional networks (Burns et al., 2014; Schroeder et al., 2020) or model parameters (Karoly et al., 2018b), that capture specific seizure properties. Using this approach, each seizure progression can be conceptualised as a pathway through the chosen feature space (Karoly et al., 2018b; Nevado-Holgado et al., 2012; Schroeder et al., 2020; Wendling et al., 2002). Second, each seizure is also characterised by its duration, which is commonly defined as the amount of time that elapses from its electrographic start to finish (Dobesberger et al., 2015; Halford et al., 2015; Kaufmann et al., 2020; Kim et al., 2011). Together, these features describe both the sequence of brain activity during a seizure as well as the amount of time that it takes to complete that sequence.

Importantly, seizure pathways and seizure durations are related to clinical symptoms (Dobesberger et al., 2015; Kaufmann et al., 2020; Kim et al., 2011; Schroeder et al., 2020), seizure severity (Cramer and French, 2001), and the outcome of treatments such as surgical resection (Burns et al., 2014) and seizure prediction (Cook et al., 2016). However, despite their clinical relevance, little is known about how these features interact. In particular, it is unclear whether variability in seizure duration arises purely from changes in seizure pathways, or whether pathways and durations can vary independently within the same patient.

Some past studies suggest that seizure pathways and durations are linked, with different durations corresponding to seizures with distinct pathways. First, seizure duration often differs between different International League Against Epilepsy (ILAE) clinical seizure types, which classify seizures based on features such as clinical symptoms (Fisher et al., 2017) and are also associated with changes in functional networks (Burns et al., 2014; Schindler et al., 2007b; Schroeder et al., 2020). Seizure types with more severe clinical manifestations tend to last longer; for example, focal seizures that progress to bilateral tonic-clonic seizures tend to have longer durations than seizures that remain focal (Dobes-

berger et al., 2015; Kaufmann et al., 2020), and focal seizures tend to be longer if they involve loss of awareness (Dobesberger et al., 2015; Kim et al., 2011). Meanwhile, analysis of chronic intracranial EEG (iEEG) recordings suggests that seizures with different durations have similar onsets, but different terminations (Karoly et al., 2018b). Additionally, there is evidence that distinct populations of short and long seizures correspond to different seizure pathways with characteristic durations (Cook et al., 2016; Karoly et al., 2018b). Overall, such findings suggest that seizure pathways and durations may co-vary within patients, with different seizure durations serving as a proxy for different seizure pathways.

However, it is also possible that seizure duration is modulated independently of seizure pathways. At one extreme, two seizures could have the same pathway, but different durations due to variability in their rates of progression. In a rodent model, Wenzel et al. (2017) found seizures with consistent recruitment patterns and different rates of seizure spread at a neuronal level, a characteristic that they termed "elasticity." To our knowledge, no studies have extensively explored such temporal flexibility in seizure pathways in human patients. Nonetheless, within-patient seizures with consistent firing patterns, but small changes in duration, have been observed (Truccolo et al., 2011), suggesting that elasticity in the same seizure pathway may also occur in humans. This mechanism could potentially lead to variable durations among seizures with the same pathway.

Thus, it is unclear whether seizure durations and seizure pathways are co-modulated within individual patients with focal epilepsy. In the past, the relationship between these features has been difficult to investigate due to the lack of an objective measure for comparing seizure pathways. In Chapter 4, we addressed this need by proposing an approach for quantitatively comparing within-patient seizure pathways and specifically investigated variability in seizure functional network evolutions (Schroeder et al., 2020). In this chapter, we used the same approach to explore if variability in seizure pathways was linked to variability in seizure durations. Importantly, our comparison of seizure pathways allowed us to recognise similar pathways even if they progressed at different rates; thus, we could determine if two seizures shared the same pathway, even if their durations differed. Moreover, we extended our analysis from the original epilepsy monitoring unit (EMU)

patients to include long-term recordings from NeuroVista patients and canines, allowing us to analyse the relationship between pathways and durations in subjects with a higher number of recorded seizures that occurred over longer timescales.

5.2 Results

To explore the relationship between seizure pathways and seizure durations, we analysed a total of 3,224 seizures, recorded using iEEG (Fig. 5.1A), from 44 subjects with focal epilepsy or focal onset seizures: 31 EMU patients (average 16.5 seizures/patient), 10 NeuroVista patients with long-term recordings (Cook et al., 2013) (average 252.6 seizures/patient) and 3 canine subjects (Howbert et al., 2014) (average 62.3 seizures/subject). Fig. 5.1B shows the iEEG recordings of four seizures from an example subject, EMU 821. In the remainder of this chapter, we first describe our approach to comparing seizure pathways and durations within individual subjects. We next use our measures to determine how pathways and duration can co-vary and independently vary within a given subject.

5.2.1 Quantifying within-subject variability in seizure pathways and seizure durations

In this work, we describe the dynamics of each seizure using two features:

1. The seizure's functional network evolution, which can be considered a *pathway* through the space of possible functional network interactions (Fig. 5.1C).

2. The amount of time that it takes the seizure to traverse its pathway; i.e., the seizure's *duration* in time (Fig. 5.1D).

For clarity, throughout this chapter we will only use the terms *short/long* to describe seizure temporal duration and *small/large* to describe relative amounts of spatial distances traversed by seizure pathways. These distinctions are important since a seizure's relative size in time and space are not necessarily linked; for example, a short seizure can have a large pathway if it progresses quickly through network space. Likewise, a long seizure with network interactions that barely change from onset to termination would only traverse a small part of network space.

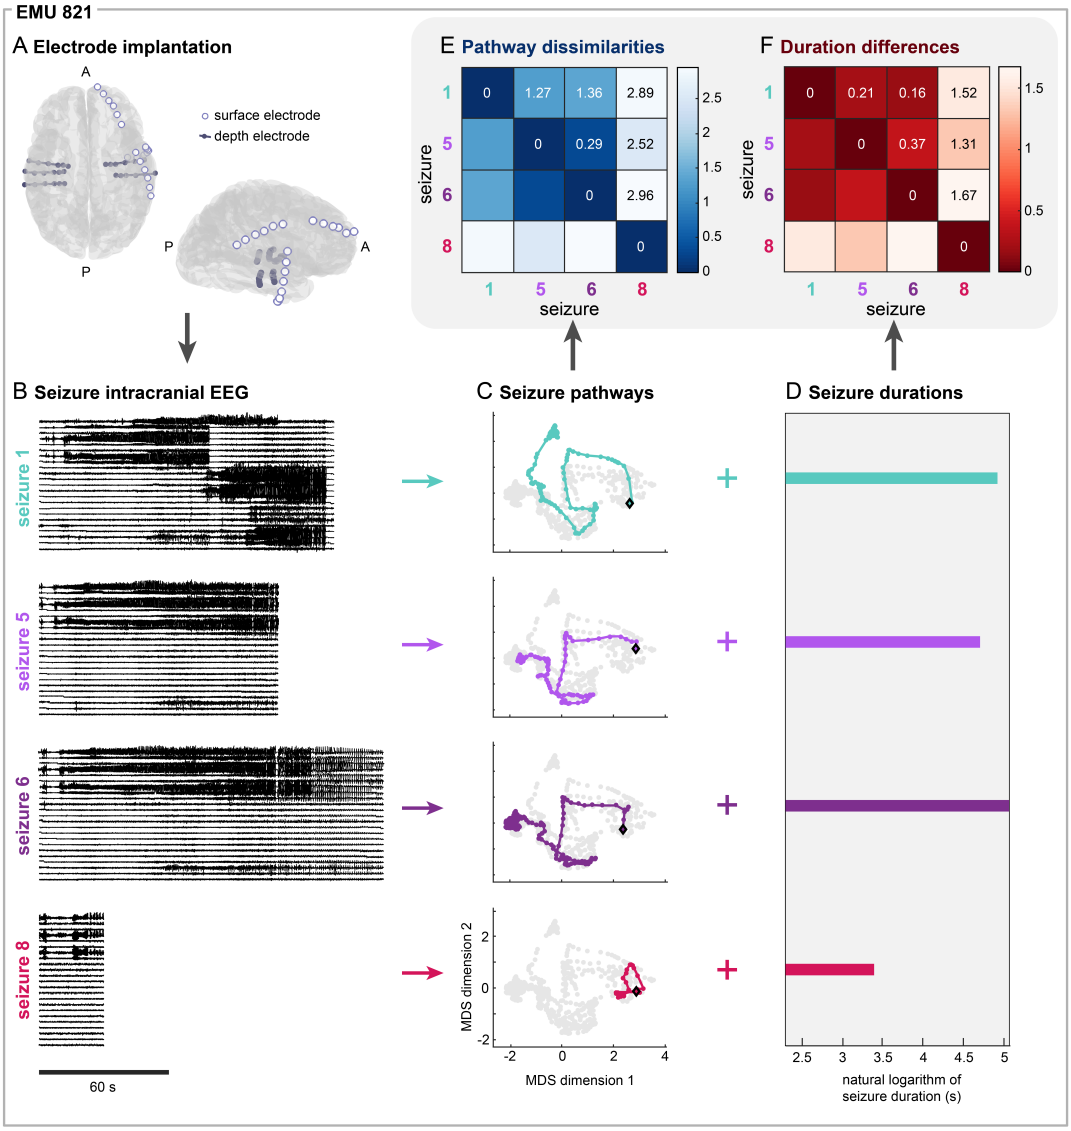

Figure 5.1: **Quantitatively comparing seizure pathways and durations within individual subjects.** A) Schematic of the electrode implantation for iEEG recording of an example subject, EMU 821. A = anterior of brain, P = posterior of brain. B) Intracranial electroencephalographic (EEG) of four of EMU 821's seizures. The recordings from a representative subset of electrodes are shown. C) Multidimensional scaling (MDS) embeddings of the corresponding pathways of the example seizures. Each point corresponds to the functional network configuration of a seizure time window, and time windows with more similar network configurations are located closer together in the embedding. Coloured points correspond to time windows that occurred during the example seizure, with the first time window marked with a black diamond and successive time window connected with the coloured line to form the seizure pathway. Time windows that occurred during other seizures are shown in grey for reference. D) The durations of each of the example seizures, shown on a natural logarithm scale. Seizure dynamics were characterised by seizure pathways (C) and seizure durations (D). E) Pairwise pathway dissimilarities and F) duration differences of the example seizures. Both matrices are symmetric.

As in previous work (Schroeder et al., 2020), we described seizure pathways through network space by computing the time-varying (sliding window) coherence between pairs of iEEG channels across six frequency bands: δ (1-4 Hz), θ (4-8 Hz), α (8-13 Hz), β (13-30 Hz), γ (30-80 Hz), and high γ (80-150 Hz). Each seizure was therefore described by a multivariate time series that captured how interactions between channels changed over the course of the seizure. To visualise seizure pathways through network space, we used multidimensional scaling (MDS) to project each subject's seizure network evolutions into a two-dimensional space (Fig. 5.1C). Each point in the projection corresponds to a network configuration that occurred during the seizure, and time points with more similar network configurations tend to be plotted closer together. These projections therefore allow us to qualitatively compare seizure pathways within the same subject. For example, in EMU 821, seizures 5 and 6 appear to have highly similar pathways relative to the other example seizures.

To quantify similarities and differences in seizure pathways, we then used dynamic time warping (DTW) (Sakoe and Seibi, 1978) to compute pairwise dissimilarities between seizures, resulting in a symmetric "seizure dissimilarity" matrix for each subject (Schroeder et al., 2020) (Fig. 5.1E). In our case, DTW minimised the overall distance between a pair of seizure pathways by selectively stretching parts of each pathway such that similar network configurations were temporally aligned. Therefore, DTW allowed us to recognise similar seizure pathways even if the seizures had different durations. We defined the *pathway dissimilarity* between a pair of seizures as the average distance between their functional connectivity time series after DTW.

To compare variability in seizure pathways to variability in seizure durations, we additionally quantified the pairwise differences in each subject's seizure durations. As in previous work (Cook et al., 2016), we first computed the natural logarithm of each seizure duration (Fig. 5.1D). We then computed the pairwise absolute differences between the transformed seizure durations, resulting in a symmetric "duration difference" matrix for each subject (Fig. 5.1F). Note that due to the properties of logarithms, our measure captures relative changes in duration (see Methods, section 5.4.2).

Thus, each subject's spatiotemporal seizure variability was described by two matrices:

their pathway dissimilarity matrix (Fig. 5.1E), containing pairwise comparisons of seizure pathways through network space, and their duration difference matrix (Fig. 5.1F), composed of pairwise differences of seizure durations. In our downstream analyses, we used these two measures to explore the relationship between seizure pathways and seizure durations in each subject. As such, our analysis focused on *differences* in seizure pathways and durations between pairs of seizures, rather than the pathway and duration features themselves. This seizure pair approach had two main advantages. First, unlike seizure duration, seizure pathways do not map onto a single feature that changes from seizure to seizure (Schroeder et al., 2020). However, we could readily measure changes in pathways with our pathway dissimilarity measure, and we could likewise compare seizure duration using duration differences. Therefore, our pairwise measures naturally allowed us to ask questions such as, "Does a pair of seizures have similar pathways if and only if they have similar durations?" Second, comparing these features at the seizure pair level was a more appropriate analysis for features that vary on a spectrum. In many subjects, seizures cannot be clearly grouped based on their pathways (Schroeder et al., 2020) or durations (Cook et al., 2016; Karoly et al., 2018b) because these features vary continuously, producing a spectrum of seizure dynamics. Forcing seizures into distinct groupings of "similar pathways" and "similar durations" would therefore have obscured the different, continuous levels of variability in these subjects. While we also analysed the special case of subjects with distinct seizure duration populations (Results section 5.2.5), our pairwise approach allowed us to precisely compare the spectrum of seizure pathways and durations present in most subjects.

5.2.2 In most subjects, seizure durations and seizure pathways are only weakly to moderately related

We first compared each subject's pathway dissimilarity matrix to their duration difference matrix. Fig. 5.2 shows the matrices of three example subjects, one from each cohort. Visually comparing the matrices within each subject already revealed that their concordance varied across subjects. Specifically, NeuroVista 11's pathway dissimilarity (Fig. 5.2A) and duration difference (Fig. 5.2B) matrices had very similar structures, with relatively low and high values in similar places in both matrices. As such, seizures with dissimilar

pathways also appeared to have different durations in this subject. On the other hand, Canine 3's matrices (Fig. 5.2G,H) had drastically different structures, suggesting little or no relationship between pathway dissimilarity and duration differences. Finally, EMU 1200's matrices (Fig. 5.2D,E) appeared to fall between these two extremes: while there were visual similarities across the two matrices, each matrix also had unique patterns of relatively low, intermediate, and high values.

To quantitatively compare these matrices, we correlated each subject's pathway dissimilarities and duration differences. This analysis captured whether seizures with similar pathways also tended to have similar durations. As expected based on the visual comparison, the example subjects in Fig. 5.2 had a range of correlations: NeuroVista 11 had a high positive correlation (Fig. 5.2C), EMU 1200 had a weak positive correlation (Fig. 5.2F), and Canine 3 (Fig. 5.2I) had approximately no correlation between pathway dissimilarities and duration differences. These examples demonstrate that the relationship between seizure pathways and seizure durations differed across subjects. In most subjects, pathway dissimilarities and duration differences were weakly to moderately correlated (Fig. 5.2J), as in EMU 1200 (median correlation: 0.322, first quartile correlation: 0.191, third quartile correlation: 0.537).

Using permutation tests, we determined that the correlations were statistically significant in 17/31 EMU patients, 9/10 NeuroVista patients, and 2/3 canine subjects after correction for multiple comparisons (Supplementary 5.5.1). The strength of the correlation between pathway dissimilarities and duration differences was not significantly associated with the maximum pathway dissimilarity or duration difference of each subject (Supplementary 5.5.2). Thus, the relationship between pathways and durations was not contingent upon or driven by the amount of variability in either feature.

Overall, the weak to moderate correlations revealed that although seizures with more similar pathways tended to also have more similar durations, changes in seizure durations were not completely explained by changes in seizure pathways and vice versa. Therefore, seizure pathways and durations contained complementary information about the dynamics of a given seizure.

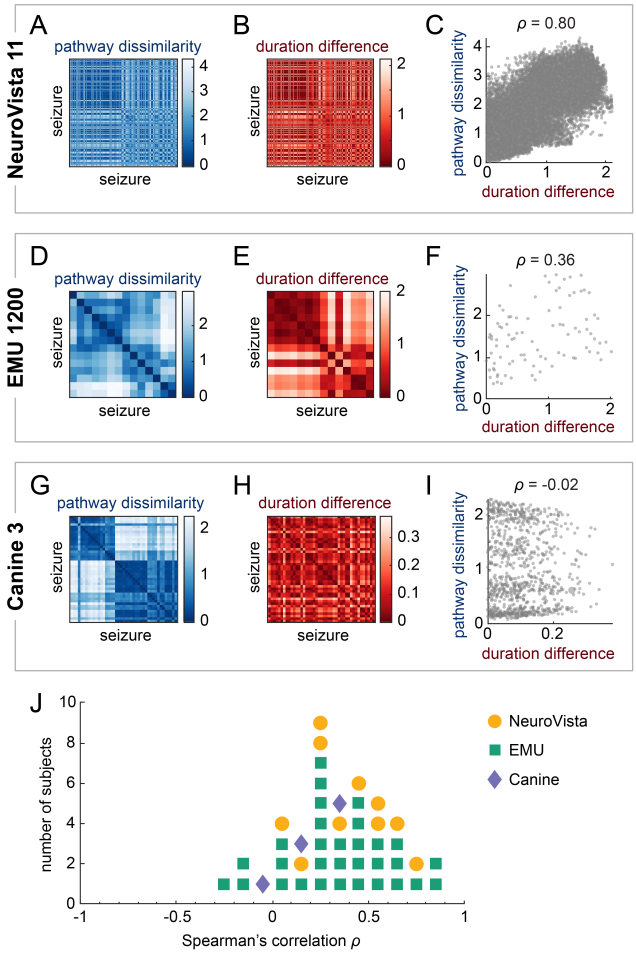

Figure 5.2: **Comparison of pathway dissimilarities and duration differences.** A-I) Comparison of pathway dissimilarities and duration differences in three example subjects. A,D,G) Pathway dissimilarity matrices of the example subjects. Each matrix quantifies the pairwise dissimilarities of the subject's seizure pathways. B,E,H) Duration difference matrices in the same subjects. Each matrix quantifies the pairwise differences in the subject's seizure durations on a natural logarithm scale. C,F,I) Scatter plots and Spearman's correlations of each subject's pathway dissimilarities vs. duration differences. Each point corresponds to a seizure pair. J) Dot plot of the Spearman's correlations between pathway dissimilarities and duration differences of all subjects. Each marker corresponds to a subject, with the colour and shape indicating the subject's cohort.

5.2.3 The relationship between pathways and durations is strengthened by pairs of seizures with both similar, or both dissimilar, pathways and durations

We next examined how pairwise relationships between seizures could strengthen or weaken the association between seizure pathways and seizure duration within each subject. A pair of seizures could fall into one of four possible categories:

1. The seizure pair had similar pathways and similar durations.
2. The seizure pair had different pathways and different durations.
3. The seizure pair had similar pathways, but different durations.
4. The seizure pair had different pathways, but similar durations.

We first found cases in which the seizure pair's pathway and duration agreement was concordant (i.e., both features similar or both features different, cases 1 and 2). Fig. 5.3 shows example pairs of seizures that had similar pathways and similar durations (case 1, Fig. 5.3A,D,G) or different pathways and different durations (case 2, Fig. 5.3B,E,H). In the latter case, we observed that the different pathways could either partially overlap in network space (Fig. 5.3B) or occupy distinct regions (Fig. 5.3E,H). Therefore, these disparate pathways could either share some common network features or have completely unrelated network evolutions.

Fig. 5.3C,F,I, visualises how these pairs of seizures impact the relationship between pathway and duration variability in each of the example subjects. When their pathways and durations were similar, the seizure pair had a low pathway dissimilarity and a low duration difference (purple points). Meanwhile, the pairs of seizures with different pathways and durations had high pathway dissimilarities and high duration differences (dark grey points). The combination of such seizure pairs within the same subject contributed to the observed positive correlations between pathway dissimilarities and duration differences. Although not visualised here, coinciding intermediate levels of changes in both seizure pathways and

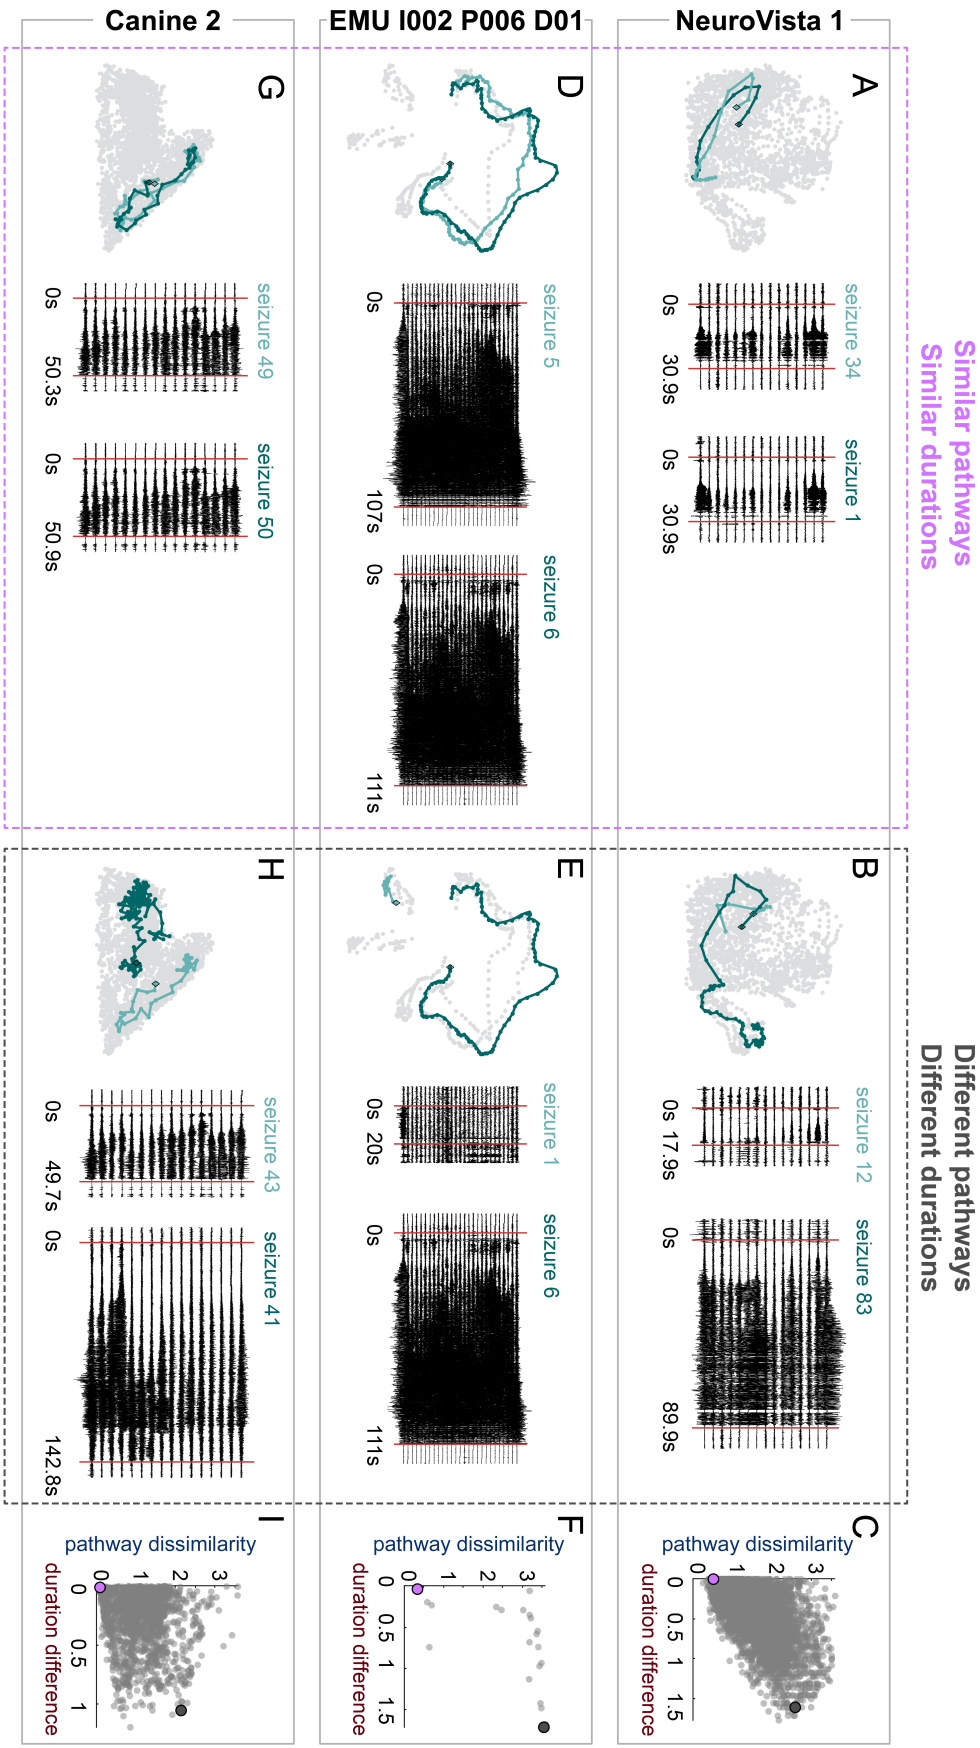
NeuroVista 1
EMU I002 P006 D01
Canine 2
Similar pathways
Similar durations
Different pathways
Different durations
A
B
C
D
E
F
G
H
I
seizure 34
seizure 1
seizure 12
seizure 83
seizure 5
seizure 6
seizure 1
seizure 6
seizure 49
seizure 50
seizure 43
seizure 41
0s
30.9s
17.9s
89.9s
20s
107s
111s
50.3s
50.9s
49.7s
142.8s
pathway dissimilarity
duration difference

Figure 5.3: **Example seizure pairs that strengthen the relationship between seizure pathways and seizure durations** (previous page). A,D,G) Example pairs of seizures with similar pathways and similar durations. Left: the seizures' pathways (light teal and dark teal), embedded in network space as in Fig. 5.1C. The time points in the subjects other seizure pathways are shown for reference (light grey points). Right: The iEEG and durations of each pair of seizures, with 10s of preictal and postictal data also shown. Red lines mark seizure onset and termination. For EMU I002 P006 D01, a representative subset of channels is shown. B,E,H) For the same subjects as in A,D, and G, example pairs of seizures with different pathways and different durations. Visualisation formats are the same as in A,D, and G. For A,B,D,E,G, and H, the time and voltage scales of the iEEG traces are consistent for each subject, but not across subjects. C,F,I) Scatter plots of pathway dissimilarities vs. duration differences of the three example subjects, with the example seizure pairs highlighted in purple (similar pathways, similar durations) and dark grey (different pathways, different durations).

seizure durations also strengthened this relationship. In other words, pathway and duration variability were related when changes in pathways produced proportional changes in durations and vice versa.

5.2.4 The relationship between pathways and durations is weakened by elastic pathways and duplicate durations

We next examined how pairs of seizures could weaken the relationship between seizure pathways and durations. First, as described in the previous section, a pair of seizure could have similar pathways, but different durations (case 3). Fig 5.4A-F provides three examples of this scenario in our cohorts. Although the seizures in each pair followed similar routes through network space, they spent different amounts of time traversing that pathway. Following previous literature (Wenzel et al., 2017), we refer to this temporal variability as *elasticity* of the seizure's pathway. Interestingly, in these examples, the pathways were not uniformly elastic; instead, there appeared to be pathway-specific locations where a pathway dwelled for different amounts of time. For example, the NeuroVista 6, seizure 5 spent relatively more time in the middle and end of the pathway (Fig. 5.4A). Due to their shared pathways and different durations, such pairs of seizures had low pathway dissimilarity and high duration differences (Fig. 5.4B,D,F blue points). Because the duration changes were not accompanied by corresponding changes in pathways, these pairs of seizures weakened the relationship between pathways and durations. These

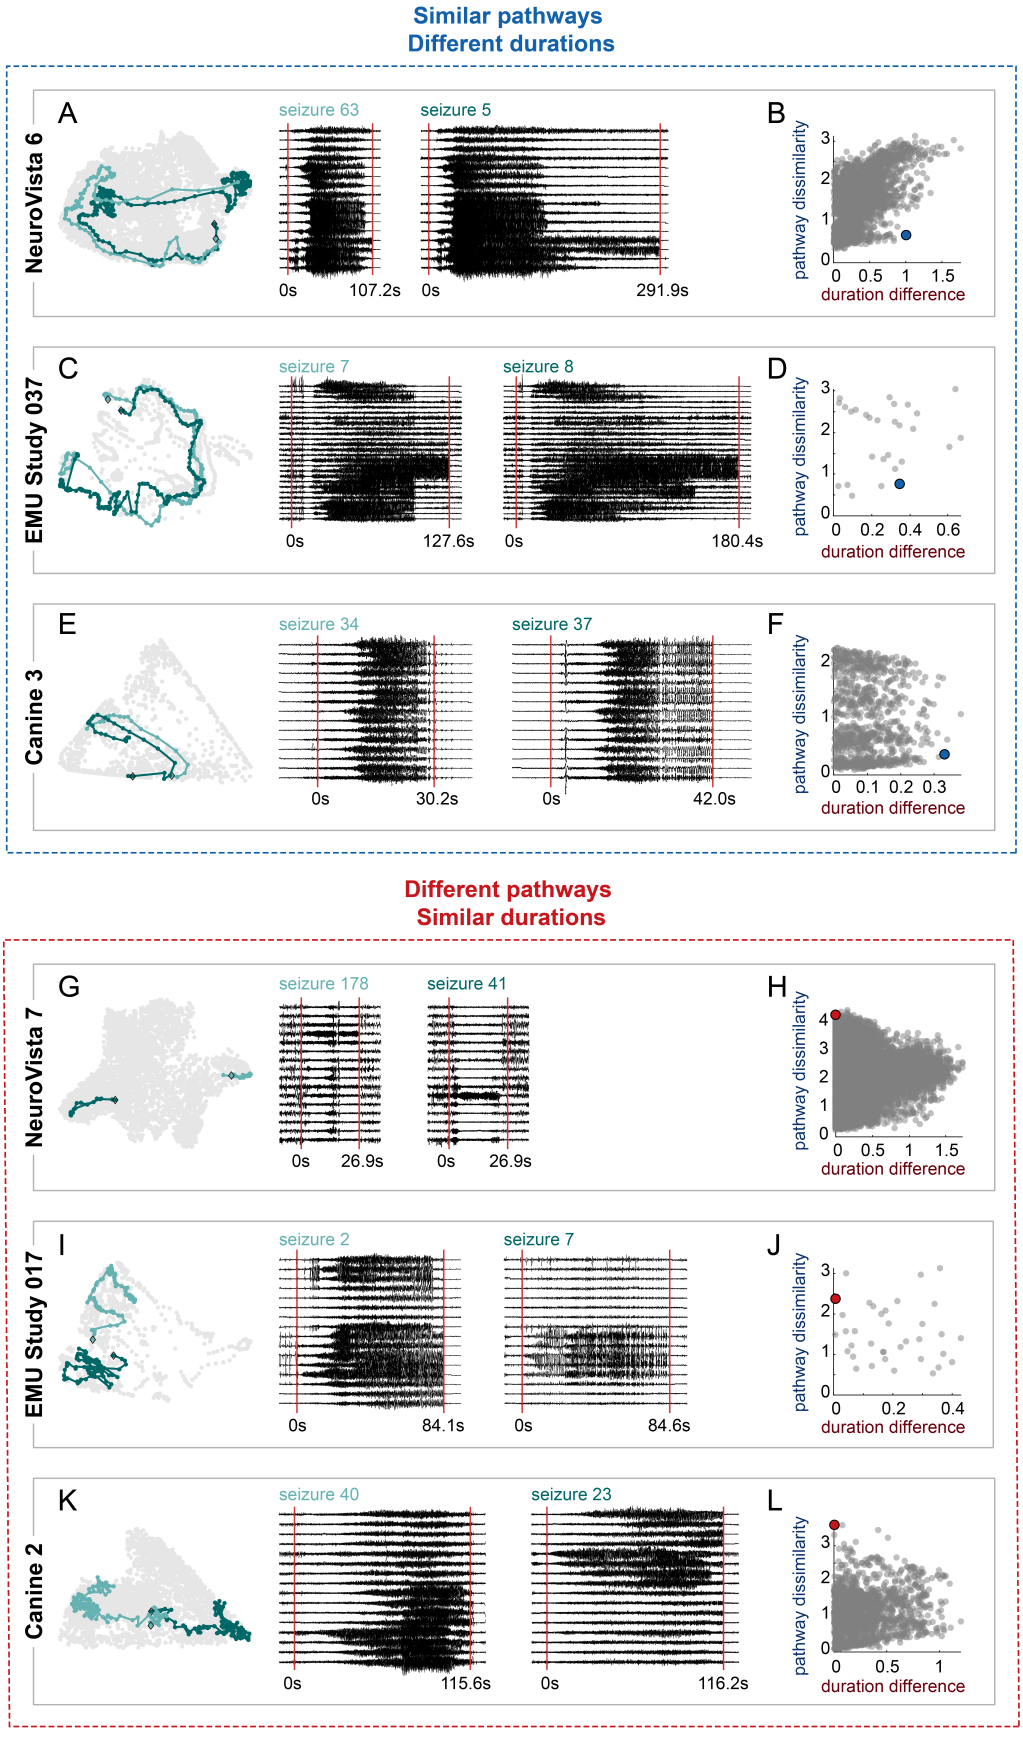
Similar pathways
Different durations
NeuroVista 6
A
seizure 63
seizure 5
0s
107.2s
0s
291.9s
B
pathway dissimilarity
duration difference
EMU Study 037
C
seizure 7
seizure 8
0s
127.6s
0s
180.4s
D
pathway dissimilarity
duration difference
Canine 3
E
seizure 34
seizure 37
0s
30.2s
0s
42.0s
F
pathway dissimilarity
duration difference
Different pathways
Similar durations
NeuroVista 7
G
seizure 178
seizure 41
0s
26.9s
0s
26.9s
H
pathway dissimilarity
duration difference
EMU Study 017
I
seizure 2
seizure 7
0s
84.1s
0s
84.6s
J
pathway dissimilarity
duration difference
Canine 2
K
seizure 40
seizure 23
0s
115.6s
0s
116.2s
L
pathway dissimilarity
duration difference

Figure 5.4: **Example seizure pairs that weaken the relationship between seizure pathways and seizure durations** (previous page). A-F) Examples of seizure pairs with similar pathways and different durations ("elastic pathways"). A,C,E) Visualisation of the seizure pathways, durations and iEEG is the same as in Fig. 5.3. For EMU Study 037, a representative subset of channels is shown. B,D,F) Scatter plots of pathway dissimilarities vs. duration differences for each subject, with the example seizure pairs highlighted in blue. G-L) Examples of seizure pairs with different pathways and similar durations ("duplicate durations"). G,I,K) Visualisation of the seizure pathways, durations and iEEG, as in Fig. 5.3. H,J,L) Scatter plots of pathway dissimilarities vs. duration differences for each subject, with the example seizure pairs highlighted in red.

results revealed that seizure durations could change independently of changes in seizure pathways. In other words, a seizure's duration was not rigidly constrained by its pathway.

The final possible scenario was that two seizures had different pathways, but the same duration (case 4). Fig 5.4G-L shows examples of this case, which we termed *duplicate durations*. These examples revealed that the duration of a seizure does not necessarily provide information about a seizure's pathway; in each of these examples, the seizures had near-identical durations, but drastically different pathways. These pairs of seizures all had low duration differences and high pathway dissimilarities (Fig. 5.4H,J,L), again weakening the relationship between pathway and duration variability in each of these subjects.

To determine the prevalence of elastic pathways and duplicate durations, we set thresholds for whether two seizures had similar pathways and/or similar durations (see Methods, section 5.4.4). We found that almost all subjects had elastic pathways (30/31 EMU, 10/10 NeuroVista, and 3/3 canine subjects) and duplicate durations (27/31 EMU, 10/10 NeuroVista, and 3/3 canine subjects) (Supplementary 5.5.3). Therefore, these mechanisms for independent variability in pathways and durations were widespread in our cohorts.

5.2.5 Populations of short and long seizures do not reliably correspond to different seizure pathways

In the previous sections, we analysed the relationship between seizure pathways and seizure durations in all subjects, regardless of the nature of their seizure dynamics. It is possible, however, that pathways and durations are more closely related in subjects

whose seizures can be grouped into different populations based on their durations. In particular, previous studies have hypothesised that patients with distinct groups of short and long seizures have different, corresponding seizure pathways (Cook et al., 2016; Karoly et al., 2018b).

As in previous work (Cook et al., 2016; Karoly et al., 2018b), we clustered seizure durations into "duration populations" in each subject and used the gap statistic to find each subject's optimal number of duration populations (see Methods section 5.4.5). While most subjects did not have multiple duration populations, a total of eight subjects (5/31 EMU patients, 3/10 NeuroVista patients, and 0/3 canine subjects) had an optimal number of two duration populations. Fig. 5.5A-F explores the pathways of these short and long duration populations in two example subjects, NeuroVista 3 and NeuroVista 8. In NeuroVista 3, there were usually relatively low pathway dissimilarities between seizures from the same duration population (i.e., between seizures that were both short or both long), but not between short and long seizures (Fig. 5.5B). These distributions indicate that the duration populations corresponded to different seizure pathways, as seen in Fig. 5.5C. Although there was still some pathway variability within each duration population (e.g., see some higher dissimilarities between long seizures, Fig. 5.5B), pathways within the same duration population were almost always much more similar than pathways from different duration populations.

Meanwhile, in NeuroVista 8 there was extensive overlap in the pathway dissimilarity distributions of within-population pairs and between-population seizure pairs (Fig. 5.5E). In other words, pairs of seizures with different durations often had more similar pathways than pairs of seizures with similar durations. As such, short and long seizures did not distinguish different seizure pathways in this patient. Seizures with similar durations could occupy drastically different parts of network space, while seizures with different durations (for example, short seizure 407 and long seizure 56) could partially overlap in network space (Fig. 5.5F). Due to the within-population pathway variability and across-population pathway similarity, seizure duration populations did not distinguish different seizure pathways in NeuroVista 8.

To quantify the agreement between seizure pathways and durations populations, we ad-

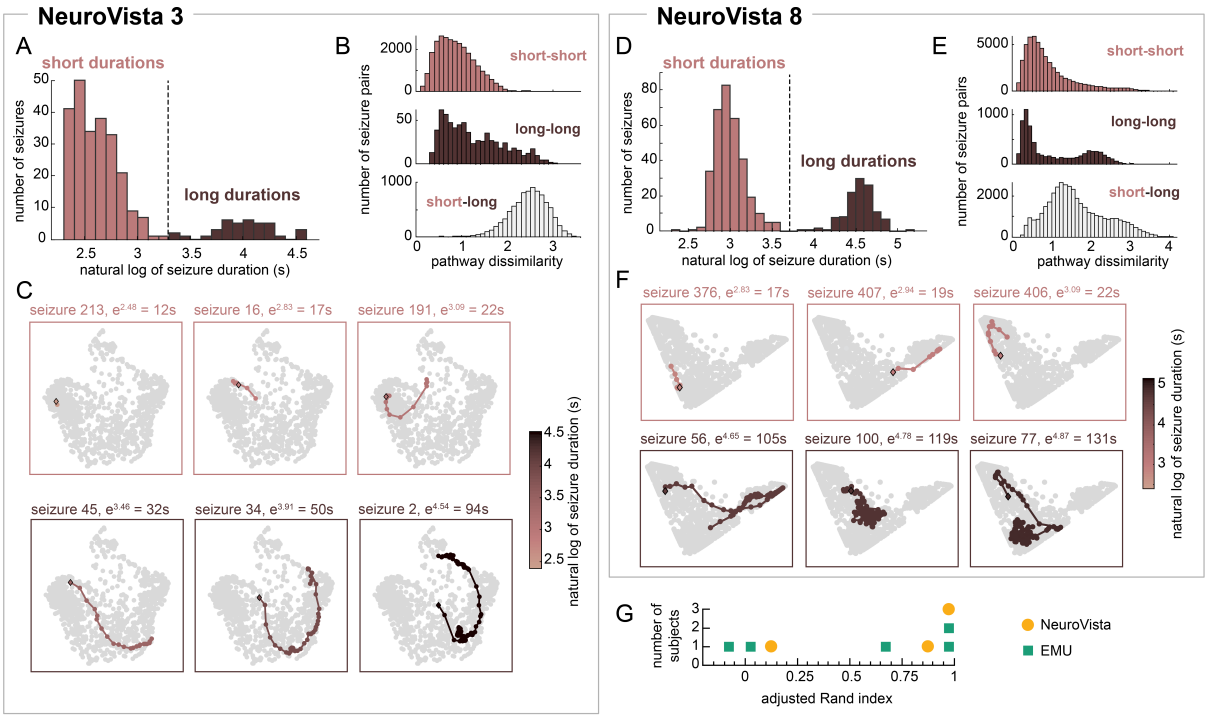

Figure 5.5: **Short and long seizures do not necessarily correspond to different seizure pathways.** Analysis of seizure pathways within subjects with multiple duration populations. A,D) Distribution of seizure durations in example subjects NeuroVista 3 (A) and NeuroVista 8 (D). Each bimodal distribution can be divided into two duration populations: one with short seizures, and one with long seizures. B,E) The distributions of pathway dissimilarities between short-short (top), long-long (middle), and short-long (bottom) pairs of seizures in each example subject. C,F) Example seizure pathways of short and long seizures in each subject. G) The adjusted Rand index between seizure duration populations and and seizure pathway clusters in all subjects with duration populations.

ditionally clustered seizures into two groups based on seizure dissimilarities. We then compared these pathway group assignments to duration populations using the adjusted Rand index (ARI) (Fig. 5.5G). An ARI of one indicated perfect agreement between the two partitions, while an ARI close to zero corresponded to only chance levels of agreement. NeuroVista 3 was one of three subjects with an ARI of one, indicating that short and long seizures perfectly corresponded to the division of seizure pathways. Meanwhile, NeuroVista 8's ARI was only 0.14; in this subject, as well as two others with near zero ARI, short and long seizures were not proxies for different seizure pathways. The remaining two subjects had intermediate levels of agreement between pathway groups and duration populations. An alternative measure of partition agreement and the statistical significance of each partition comparison are available in Supplementary 5.5.4.

Overall, these results revealed a complex, subject-specific relationship between seizure durations and pathways in subjects with multiple duration populations. We found that in some subjects, duration populations indeed corresponded to different seizure pathways, although there was still additional pathway variability within each duration population. However, in many other subjects, duration populations were not associated with different groups of seizure pathways. As such, in the absence of information about a subject's seizure pathways, duration populations are an unreliable way to distinguish different seizure pathways.

5.2.6 The relationship between seizure pathways and seizure durations is not associated with surgical outcome or seizure localisation

In the EMU patients, we related measures of seizure pathway and duration variability to the following clinical measures: ILAE surgical outcome, seizure localisation (temporal vs. extratemporal and left vs. right hemisphere seizure onset), and disease duration (see Methods, section 5.4.6). There were no significant relationships between seizure pathway/duration measures and surgical outcome after false discovery rate (FDR) correction for multiple comparisons, although there was a trend (Spearman's correlation $\rho = 0.45, p = 0.0209$) for surgical outcome to worsen as median seizure duration increased. There were also no significant differences between patients with different onset

locations, and the only trend was for patients with temporal onset to have higher median seizure duration than patients with extratemporal onset (Wilcoxon rank sum test, $p = 0.0381$).

Disease duration was significantly correlated to median seizure duration (Spearman's correlation $\rho = -0.56, p = 0.0011$), with median seizure duration decreasing as disease duration increased. Additionally, the tendency of of similar seizure pathways to be elastic significantly decreased with disease duration (Spearman's correlation $\rho = -0.51, p = 0.0038$). Importantly, disease duration at the time of the iEEG recording was not sampled at a random time for each patient, but rather determined by the clinical decision to undergo presurgical monitoring. Therefore, disease duration in our cohort could also be associated with clinical considerations such as seizure severity and the patient's level of antiepileptic medication resistance. As such factors could also influence measures such as seizure duration, it is difficult to interpret the observed associations with disease duration.

5.3 Discussion

In this chapter, we quantitatively compared two features of seizure dynamics: seizure durations and seizure pathways. We found that although these features were linked, they often varied independently within individual subjects. Seizures with the same pathway could have different durations due to temporally elastic progressions, and seizures with the same duration could have different pathways. Additionally, we found that the level of association between pathways and durations was subject-specific: in some subjects, similar seizures usually had similar durations, while in others, there was little to no association between pathway and duration similarity. Even in subjects with distinct populations of short and long seizures, seizures with different durations did not necessarily correspond to different seizure pathways. Our results indicate that seizure pathways and durations carry complementary information about a seizure's dynamics, and that these features can perhaps be modulated independently within a given subject to create the observed diversity in both pathways and durations.

One mechanism that creates variability in seizure durations, while maintaining the same pathway, is variability in the rate of a seizure's progression. Wenzel et al. (2017) previ-

ously observed this phenomenon in a rodent model, describing it as "elasticity" of seizure propagation. To our knowledge, seizure elasticity has not previously been quantitatively described in human seizures, although human seizures with consistent spiking patterns, but slightly different durations, have been observed (Truccolo et al., 2011). Our work reveals that temporal elasticity is also a common feature of human seizures. Interestingly, it appeared that such elasticity did not necessarily affect the entire seizure pathway; instead, a seizure could selectively dwell in certain parts of a given pathway. Further research is needed to understand what parts of seizure pathways are most prone to variable rates of progression, as well as the underlying mechanisms, such as local (Wenzel et al., 2017) or feedforward (Trevelyan et al., 2007) inhibition, that determine these temporal features. Uncovering these mechanisms could provide possible clinical strategies for controlling seizure progression. For example, the same mechanism that slows seizure propagation could potentially be leveraged to entirely halt seizure progression at these dwell sites, limiting seizure spread.

It is also likely that some mechanisms affect both seizure pathways and durations, creating seizures that have both different pathways and different durations. One possibility is that temporally shorter and longer seizures can arise from truncating seizure pathways (Karoly et al., 2018b). This process could produce temporally longer seizures that progress along the entire pathway as well as temporally shorter seizures that only traverse the initial part of the full pathway. Indeed, we observed that temporally short and long seizures can have overlapping pathways (Fig. 5.3B). In support of this hypothesis, Karoly et al. (2018b) found that within-patient seizure duration was associated with seizure terminations, but not onsets. This observation suggests that the seizures with different durations shared the same initial pathway. Additionally, microelectrode recordings have revealed that some patients have shorter seizures that terminate earlier along the patient's characteristic seizure progression, again suggesting that shorter seizures can arise by truncating seizure pathways (Wagner et al., 2015). While this mechanism could be responsible for creating duration populations with different pathways, it could also exist in patients with a spectrum of seizure durations (Karoly et al., 2018b). When present in a patient, a truncation mechanism could potentially be exploited to trigger early termination of seizure

pathways and thereby reduce seizure duration and severity. Future studies could determine the prevalence of truncated seizure pathways as well as the specific pathways that are susceptible to truncation in each patient.

As most subjects had a spectrum of seizure pathways and durations, the majority of our analysis compared these measures on a seizure pair level. However, we additionally analysed temporally short and long seizures in patients with distinct duration populations, as these duration populations are thought to correspond to different seizure pathways (Cook et al., 2016; Karoly et al., 2018b). Consistent with previous analyses (Karoly et al., 2018b), we identified three NeuroVista patients with duration populations. We augmented our analysis with a small number of EMU patients with apparent duration populations, although long-term recordings of these patients would be needed to determine if these populations persist over longer time periods. Interestingly, duration populations tended to correspond to different, distinct groups of seizure pathways in two of the three NeuroVista patients, with little relationship between seizure durations and pathways in the remaining patient. Additionally, even when duration populations segregated seizure pathways, there was additional, unexplained pathway variability within each duration population. The EMU patients showed similar trends. These results indicate that while duration populations can be signs of different seizure pathways, they are not reliable proxies for seizure pathways without additional information about the patient's pathway characteristics. Further research is needed to determine mechanisms, such as pathway truncation (Karoly et al., 2018b), that could create distinct duration populations that correspond to different pathways.

The factors that could independently modulate and co-modulate seizure pathways and durations are also unknown. We have previously hypothesised that preictal variability in brain dynamics could produce changes in seizure pathways (Schroeder et al., 2020). Likewise, variability in preictal features such as levels of cortical excitability and inhibition could potentially affect seizure duration. Importantly, the observed independent variability in seizure pathways and seizure durations suggests that these features can be modulated separately in some patients. In other words, some factors may solely shape seizure spread, patterns, and connectivity, while others could determine the rate of seizure

evolution. Disentangling these effects may require accounting for variability in one feature when analysing the other aspect of seizure dynamics. For example, it may be difficult to uncover preictal associations with seizure duration without controlling for variability in seizure pathways.

Our study was limited to human patients with focal epilepsy and canines with focal-onset seizures, and it is unclear whether similar relationships between seizure pathways and durations exist in other types of epilepsies. However, our concordant findings in canine subjects indicates that our results generalise beyond human patients. It is also likely that we did not observe all types and combinations of pathway and duration variability in our subjects, especially in EMU patients with shorter recordings (King-Stephens et al., 2015); for example, it is impossible to observe elasticity in a seizure pathway that only occurs once in a recording. Another limitation of our study is that seizure duration depends on clinically or algorithmically marked seizure onsets and terminations. Clinical markings can be subjective and vary from marker to marker, especially in some seizures with more ambiguous onsets (Davis et al., 2018; Halford et al., 2015). Seizure duration errors were likely small relative to the length of most seizures, however, and were therefore unlikely to drive our results.

Overall, we have revealed that seizure pathways and durations can vary independently, increasing the possible combinations of seizure dynamics that can occur in a given patient. As such, both pathway and duration information is needed to fully characterise a seizure's dynamics. Determining the mechanisms by which each feature independently and co-varies could potentially provide strategies for slowing or terminating seizure spread in therapeutic interventions.

5.4 Methods

The steps for seizure selection and extraction, iEEG preprocessing steps, and computation and dimensionality reduction of seizure functional connectivity are available in Chapter 2. Additionally, the EMU subjects and seizure functional connectivity used in this study were the same as the cohort in Schroeder et al. (2020) (Chapters 3 and 4). These initial steps

transformed each subject's seizure iEEG data into functional connectivity (i.e., functional network) time series, which we described as pathways through network space.

5.4.1 Comparing seizure pathways using pathway dissimilarities

We used the approach of Schroeder et al. (2020) to compare pairs of seizure pathways, which were described by the functional connectivity time series, within each patient (see Chapters 2 and 3 for details). Briefly, for each pair of seizures, we used DTW (Sakoe and Seibi, 1978) (MATLAB function *dtw*) to align similar time points in their functional connectivity time series and minimise the overall $L1$ distance between the time series. The seizure pair's "pathway dissimilarity" (previously called "seizure dissimilarity" in Schroeder et al. (2020) and Chapters 3 and 4) was then defined as the average $L1$ distance between their warped time series. Repeating this process for each pair of a subject's s seizures yielded the subject's pathway dissimilarity matrix, a symmetric $s \times s$ matrix containing all of the pairwise pathway dissimilarities.

5.4.2 Comparing seizure durations using duration differences

To compare seizure durations, we computed a pairwise "duration difference" measure for each pair of a subject's seizures. First, as in previous work (Cook et al., 2016), we transformed each subject's seizure durations by computing their natural logarithm, which made each subject's distribution of seizure durations closer to a normal distribution. We then defined the duration difference between a pair of seizures as the absolute difference between their transformed durations,

$$|ln(l_i) - ln(l_j)|$$

where l_i and l_j are the durations, in seconds, of seizures i and j, respectively. Due to the properties of logarithms, this measure is equal to

$$|ln(l_i/l_j)|$$

and therefore depends on the ratio between the durations of seizures i and j. As such, duration differences capture the proportional differences between seizure durations. For

example, the duration difference between a 20s seizure and a 40s seizure will be the same as the duration difference between a 60s seizure and a 120s because in both cases, the longer seizure is twice the duration of the shorter seizure. Likewise, a certain absolute change in duration, such as 10s, results in a larger duration difference when the original seizure is shorter. As for the pathway dissimilarity measure, duration differences were computed for each pair of a subject's s seizures to create the subject's symmetric $s \times s$ duration difference matrix.

5.4.3 Comparing pathway dissimilarities and duration differences

To compare pathway dissimilarities and duration differences, we used the same approach as Schroeder et al. (2020) for comparing dissimilarity matrices. Specifically, for each subject, we computed Spearman's correlation between the upper triangular elements of their pathway dissimilarity and duration difference matrices. We then used the Mantel test (Mantel, 1967) (10,000 permutations, one-sided significance test) to determine the probability of obtaining a correlation greater than or equal to the observed correlation by chance.

5.4.4 Defining elastic pathways and duplicate durations

To determine the prevalence of elastic pathways and duplicate durations, we set first thresholds for defining whether a seizure pair had similar pathways (pathway dissimilarity ≤ 1) and similar durations (duration difference ≤ 0.2). The pathway threshold was chosen because seizures with pathway dissimilarities below the threshold tend to have visually similar pathways as well as electrographic patterns. The duration difference threshold allows a $e^{0.2} = 1.22$ fold increase in duration relative to the shorter seizure before the durations are considered difference. These thresholds were also set so that the overall proportion, across all subjects, of seizure pairs with similar pathways was comparable to the proportion of seizure pairs with similar durations (32.6% and 34.7%, respectively). Seizures pairs with similar pathways (pathway dissimilarity ≤ 1) and different durations (duration difference > 0.2) were then defined as examples of elastic pathways. Seizure pairs with different pathways (pathway dissimilarity > 1) and similar durations (duration

difference ≤ 0.2) were considered examples of duplicate durations.

5.4.5 Comparing duration populations and seizure pathways

In each subject, we determined the number of duration populations by clustering seizure durations (after the natural logarithm transformation) using k-means, with the number of clusters k scanned from 1 to 5. The gap statistic (Tibshirani et al., 2001) (MATLAB *evalclusters*, search method *firstMaxSE*, with 1000 reference distributions) was used to select the optimal number of clusters, with $k = 1$ indicating an absence of multiple duration populations and $k \geq 2$ revealing multiple duration populations.

In subjects with multiple duration populations, we additionally clustered seizures into an equivalent number of groups k to compare seizure clusters based on duration with clusters based on pathways. In each subject, seizure pathways were clustered by applying unweighted pair group method with arithmetic mean (UPGMA) hierarchical clustering to the pathway dissimilarity matrix, and the resulting dendrogram was then cut to produce the same number of discrete pathway clusters as duration populations. The Rand index and adjusted Rand index were then computed to compare the duration population and pathway clusters partitions. To determine the statistical significance of these measures, the cluster membership for one partition was permuted 10,000 times and the measures were recomputed for each partition to create null distributions.

5.4.6 Clinical metadata comparisons

To determine if patterns of pathway and duration variability were related, we compared nine subject-specific pathway and duration measures to four clinical variables in the EMU patients. The pathway and duration measures were

1. The correlation between the subject's pathway dissimilarities and duration differences.

2. The subject's maximum duration difference.

3. The subject's median duration difference.

4. The subject's longest seizure.

5. The subject's median seizure duration.

6. The proportion of the subject's seizure pairs with similar pathways that were elastic.

7. The proportion of the subject's seizure pairs with similar durations that were duplicate durations.

8. The median duration difference of the subject's elastic pathways.

9. The median pathway dissimilarity of the subject's duplicate durations.

These measure were compared to

1. The patients' ILAE surgical outcomes ($n = 26$).

2. The patients' disease durations ($n = 31$).

3. Whether the patient had temporal ($n = 12$) or extratemporal ($n = 15$) epilepsy.

4. Whether the patient had seizures with left ($n = 15$) or right ($n = 13$) hemisphere onset.

ILAE surgical outcomes and disease duration were associated with the variability measures using Spearman's correlation. The variability measures of temporal versus extratemporal patients and left versus right hemisphere onset patients were compared using Wilcoxon rank sum tests.

5.4.7 Correction for multiple comparisons

The Benjamini-Hochberg FDR correction (Benjamini and Hochberg, 1995), with $\alpha = 0.05$ ($\alpha = 0.10$ for trends), was applied to the set of p-values from all statistical tests in this chapter and corresponding supplementary material. Uncorrected p-values are reported in the text.

5.5 Supplementary

5.5.1 Significance test results for pathway dissimilarities/duration differences correlations

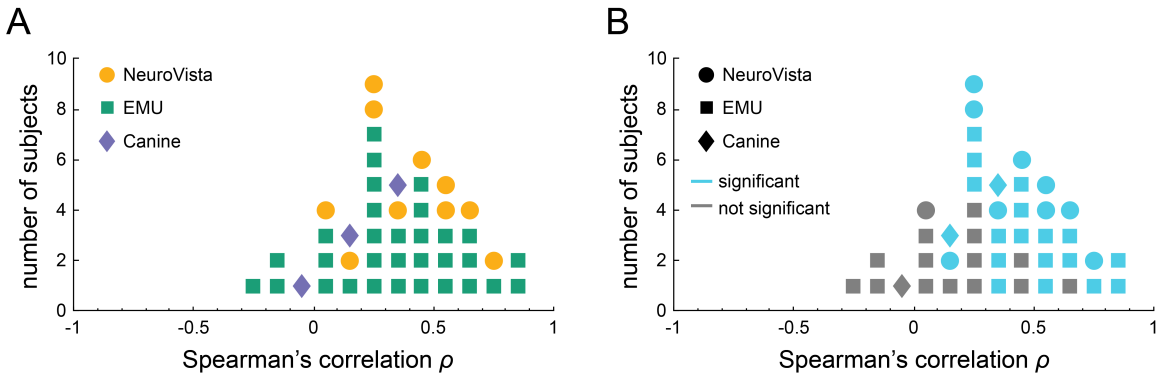

Figure 5.6: **Significance test results for pathway dissimilarities/duration differences correlations.** A) Reproduction of Fig. 5.2J showing the distribution of correlations between pathway dissimilarities and duration differences across subjects. Each marker corresponds to a subject, with the marker shape and colour indicating the subject's cohort. B) The same distribution as in A, but with markers now coloured by whether the subjects correlation was significant after FDR correction for multiple comparisons.

Fig. 5.6 shows which correlations between pathway dissimilarites and duration differences were significant after FDR correction for multiple comparisons. The Mantel test was used to test the statistical significance of each correlation (see Methods, section 5.4.7). Note that the significance of each correlation depended not only on the strength of the correlation, but also the number of the seizures in each subject and the relationships between the subject's seizure pathways. As such, some relatively high correlations may not be significant and vice versa.

5.5.2 The relationship between pathway and duration variability does not depend on the amount of variability in either feature.

It is possible that the relationship between seizure pathways and durations requires a certain amount of variability in pathways and duration. To test this hypothesis, we compared the correlation between pathway dissimilarities and duration differences to the maximum duration difference (Fig. 5.7A) and maximum pathway dissimilarity (Fig. 5.7B)

in each subject. The maximum duration difference describes the greatest proportional change in seizure duration within each subject, while the maximum pathway dissimilarity captures the largest level of variability in the subject's seizure pathways. In other words, these measures serves as ranges for each subject's seizure durations and seizure pathways.

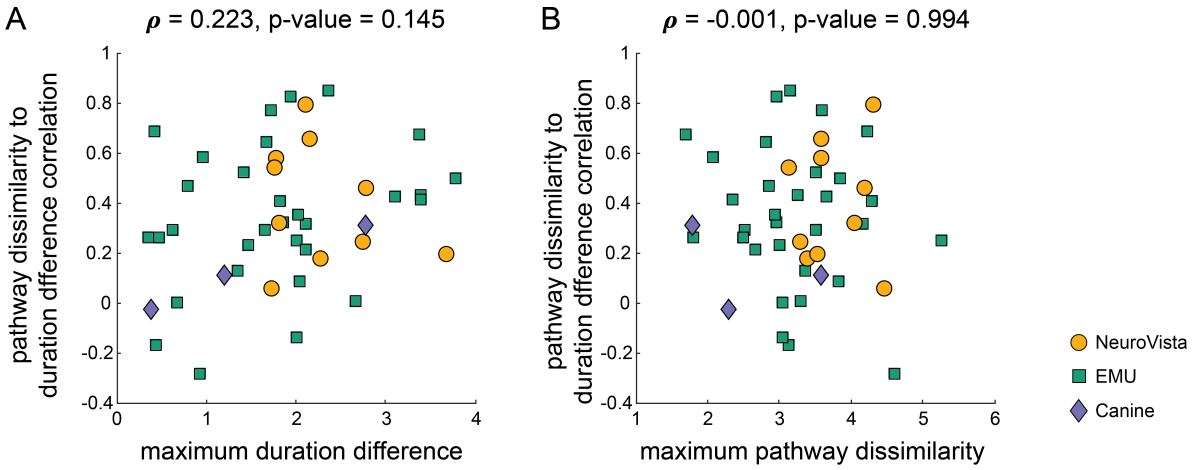

Figure 5.7: **The relationship between pathway and duration variability does not depend on the amount of variability in either feature.** A) The correlation between pathway dissimilarities and duration differences plotted versus the maximum duration difference of each subject. The colour and shape of each marker corresponds to the subject's cohort. B) The correlation between pathway dissimilarities and duration differences plotted versus the maximum pathway dissimilarity of each subject.

In both cases, we found no significant relationship between the extent of variability and the correlation between pathway dissimilarities and duration differences. There was a weak, but insignificant, positive association between pathway dissimilarity/duration difference correlations and maximum durations differences. As such, there was a slight tendency for subjects with greater duration variability to have a stronger relationship between pathways and durations. Future work could investigate this relationship in a larger cohort. However, overall, the relationship between pathways and durations does not depend on the amount of variability in these features.

5.5.3 Prevalence and features of elastic pathways and duplicate durations

Fig. 5.8 describes the prevalence of elastic pathways and duplicate durations in each subject, as well as the level of elasticity and duplication (see Methods, section 5.4.4, for

how elastic pathways and duplicate durations were defined). Fig. 5.8A demonstrates that it was common for a high proportion of seizures with similar pathways to have different durations (i.e., elastic pathways). Likewise, Fig. 5.8B reveals that in many subjects, seizures with similar durations had different pathways (i.e., they were duplicate durations).

The level of pathway elasticity, which can be quantified by the duration differences of the elastic seizure pairs, varied across subjects (Fig. 5.8C,E). On average, duration differences of 0.2 to 0.55 (equivalent to a $e^{0.2}$ to $e^{0.55} = 1.22$ to 1.73 fold change in seizure duration), were common, but extremes of $e^1 = 2.72$ or higher were also observed in many subjects. Thus, seizures with similar pathways could have drastically different durations. Likewise, the level of pathway dissimilarity between seizure pairs with similar durations varied across subjects (Fig. 5.8D,F). Average pathway dissimilarities of approximately 1.2 to 2 were common, but higher dissimilarities of 3 or higher were also observed in many subjects. Therefore, seizures with similar durations could have very different pathways.

5.5.4 Comparison of duration populations and pathway dissimilarities

Fig. 5.9 shows the adjusted Rand indices (Fig. 5.9A,B) and Rand indices (Fig. 5.9C,D) between duration populations and pathway clusters in subjects with two duration populations (see Methods). The significance of each index will depend on the index value as well as the number of seizures and cluster sizes. The Rand index is the proportion of seizure pairs that have the same relationship in both partitions (i.e., in the same cluster in both partitions or in different clusters in both partitions), and is therefore easily interpretable. However, the Rand Index greatly depends on the relative cluster sizes, making the adjusted Rand index a better measure for understanding the strength of the agreement of the two partitions. We therefore provide both measures here to evaluate the agreement between duration populations and pathway clusters.

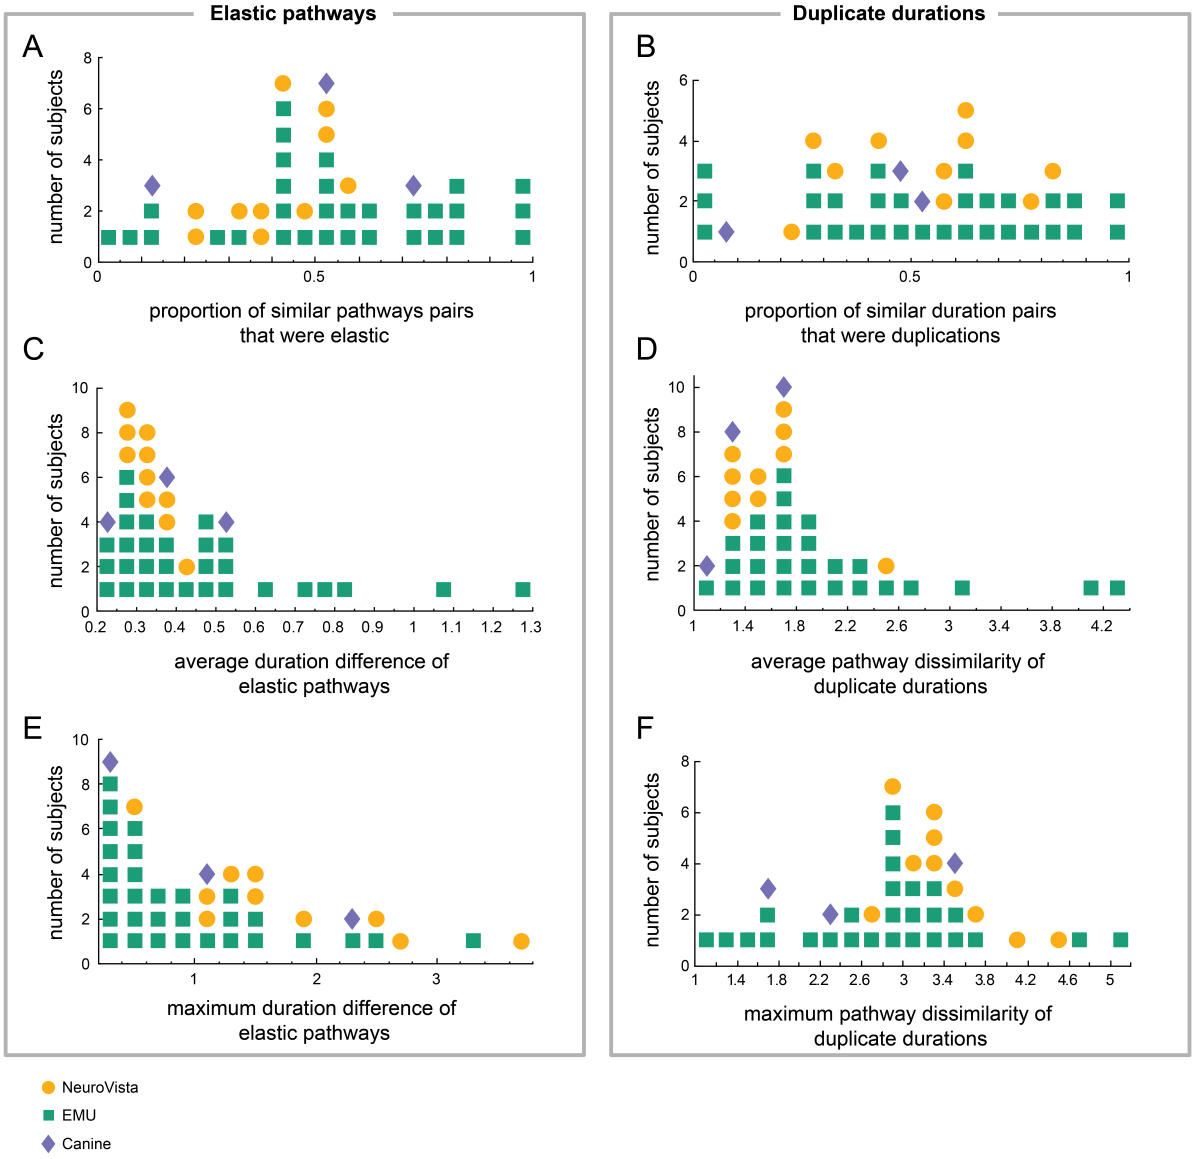

Figure 5.8: **Prevalence and features of elastic pathways and duplicate durations.** In all plots, markers correspond to subjects and their colour and shape indicates the subjects' cohorts. A) Proportion of seizure pairs with similar pathways that were elastic (i.e., that had different durations) in each subject. B) Proportion of seizure pairs with similar durations that were duplications (i.e., that had different pathways). C,E) In subjects with elastic pathway pairs, the median (C) and maximum (E) duration difference of the elastic pathways. D,F) In subjects with duplicate duration pairs, the median (D) and maximum (F) pathway dissimilarity of the duplicate durations.

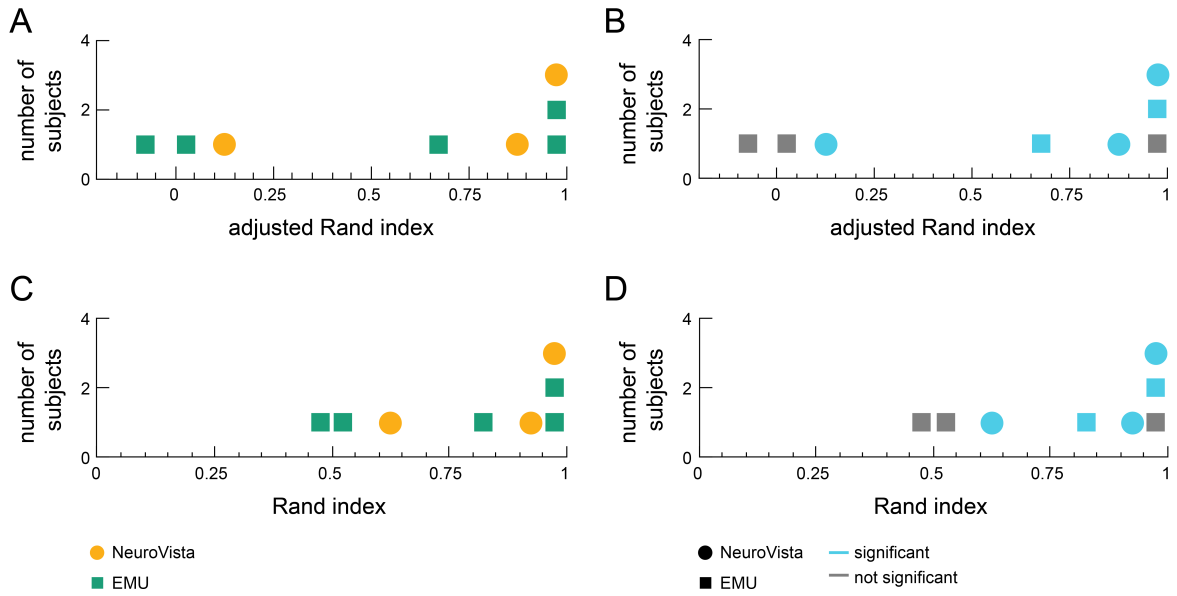

Figure 5.9: **Comparison of duration populations and pathway dissimilarities.** A) Reproduction of Fig. 5.5I showing the distribution of adjusted Rand indices between duration populations and pathway clusters. Marker colour and shape correspond to each subject's cohort. B) Same as A, but marker colour indicates whether the adjusted Rand index of the subject was significant after correction for multiple comparisons. C-D) The corresponding Rand indices of the same subjects.

Chapter 5: Seizure Pathways and Seizure Durations Can Vary Independently Within Individual Patients with Focal Epilepsy

Chapter 6. Interictal Spike Rate Reveals Timescales of Modulation in Seizure States and Seizure State Durations

Contents

6.1 Introduction

Focal epilepsy is characterised by recurrent, unprovoked seizures. Importantly, these seizures are not homogeneous events, even in the same patient. Within individual patients, seizure features such as clinical symptoms (Noachtar and Peters, 2009), onset locations and patterns (Gliske et al., 2018; Jiménez-Jiménez et al., 2015; King-Stephens et al., 2015; Saggio et al., 2020; Salami et al., 2020), durations (Cook et al., 2016; Schroeder et al., 2021) (see also Chapter 5), and network evolutions (Schroeder et al., 2020) can change over time and potentially influence treatment responses (Cook et al., 2016; Ewell et al., 2015; Ryzi et al., 2015). As such, a better understanding of the patterns and sources of within-patient seizure variability is needed.

One open question is whether and how such seizure features change over short (e.g., circadian) and long (e.g., weekly, monthly, and yearly) periods of time. There is some evidence that seizures are modulated over different timescales. For example, certain clinical seizure types and symptoms, such as secondary generalisation, can preferentially occur during specific parts of sleep/wake or day/night cycles (Bazil, 2018; Bazil and Walczak, 1997; Janz, 1962; Loddenkemper et al., 2011; Sinha et al., 2006). Electrographic seizure onset patterns can shift across the days of epilepsy monitoring unit (Gliske et al., 2018) and months of chronic iEEG (Ung et al., 2016) recordings, suggesting that seizure features can also change over slower timescales. More recently, our preliminary analysis in epilepsy monitoring unit patients found that seizure network evolutions do not change randomly over time (Schroeder et al., 2020). Instead, over the timescales of these recordings, seizures with more similar network evolutions tended to occur closer together than less similar seizures. Further, in most patients, the changes in seizure network evolutions could be explained by a combination of circadian and/or slower time-varying factors. However, these temporal associations and the specific timescales of seizure changes need to be further explored in longer recordings with larger numbers of seizures.

In recent years, chronic iEEG recordings with durations of months to years have provided unprecedented insights into epileptic brain dynamics over multiple timescales (Cook et al., 2013; Davis et al., 2011; Howbert et al., 2014; Jarosiewicz and Morrell, 2021). First, these

recordings have revealed fluctuations in interictal dynamics, including in the rates and spatial patterns of bursts (Ung et al., 2016), spikes (Chen et al., 2021), high frequency activity (Chen et al., 2021), and other signal features (Ung et al., 2017). This variability is especially high in the first months after electrode implantation, possibly due to the brain's response to acute trauma (Chen et al., 2021; Ung et al., 2017, 2016). However, more persistent variability in such features has also been observed (Chen et al., 2021; Ung et al., 2016), suggesting that other mechanisms also drive the observed interictal shifts. In addition, multiple studies have found prevalent patient-specific circadian, multidien (multi-day), and/or circannual cycles in interictal features and seizure occurrence (Baud et al., 2018; Chen et al., 2021; Karoly et al., 2018a, 2021; Leguia et al., 2021; Maturana et al., 2020). Since the exact periods of these cycles often vary over time, they are best tracked using fluctuations in continuous biomarkers such as interictal spike rate (Baud et al., 2018; Leguia et al., 2021; Maturana et al., 2020) that are thought to reflect underlying endogenous and external factors that influence brain dynamics over multiple timescales (Karoly et al., 2021; Rao et al., 2020). An intriguing possibility is that seizure features and seizure evolution could also change over such cycles. For example, certain seizure features may preferentially occur at a particular phase, while other features may gradually change over the course of a spike rate cycle. However, the relationship between seizure features and spike rate cycles has not been explored.

In this chapter, we addressed these questions by analysing changes in seizure network evolutions in chronic iEEG recordings from the NeuroVista dataset (Cook et al., 2013). We described seizure evolutions as a sequence of a small number of patient-specific network states, similar to past studies (Burns et al., 2014; Khambhati et al., 2015). In each patient, we then analysed changes in seizure states over multiple timescales. We first identified gradual changes in seizures states across the course of each recording. We then determined if seizure states also fluctuated over patient-specific circadian and multidien cycles that were revealed by interictal spike rate. To account for possible independent variability in seizure evolutions and seizure duration (Chapter 5), we separately examined variability in seizure state occurrence and seizure state duration. We show that in most patients, both of these features were associated with multiple timescales, providing new insight into the

patterns and possible mechanisms of within-patient seizure variability.

6.2 Results

We analysed seizure network evolutions in 10 patients with focal epilepsy who underwent chronic continuous iEEG recordings as part of the NeuroVista seizure prediction study (Cook et al., 2013). Patient details are provided in Chapter 2.

6.2.1 Seizure network evolutions vary from seizure to seizure within individual patients

In each patient, we first characterised changes in seizure network evolutions over the length of the patient's continuous iEEG recording. Specifically, we described seizure network evolutions as a sequence, or progression, of network states (see Methods for details). Similar approaches have previously been used to summarise seizure evolutions by characterising recurring network patterns of seizures (Burns et al., 2014; Khambhati et al., 2015). Fig. 6.1A shows these progressions for an example patient, NeuroVista 1. The vertical lines indicate seizures that occurred during the patient's approximately two year recording. The line colours indicate the network state in each seizure time window, with each time window corresponding to approximately a second of the iEEG recording. NeuroVista 1 had six seizure network states that each described a specific pattern of coherence between the patient's iEEG channels (Supplementary section 6.5.1). These state progressions therefore summarised the evolution of channel interactions during the patient's seizures. Although these states did not explicitly capture visual differences in seizure iEEG traces, seizures with different network evolutions also tend to have different electrographic patterns (Schroeder et al., 2021, 2020). As an example, Fig. 6.1B shows two seizures that began with the same three states (states B, E, and C), but had different final states (state D, dark blue, vs. state F, dark orange), at which time the amplitude and frequency of their seizure activity also diverged.

From Fig. 6.1A, it is already apparent that seizure network evolutions differed from seizure to seizure. In particular, there was variability in both state occurrence (whether a state occurred during a seizure) (Fig. 6.1C) and state duration (the total number of

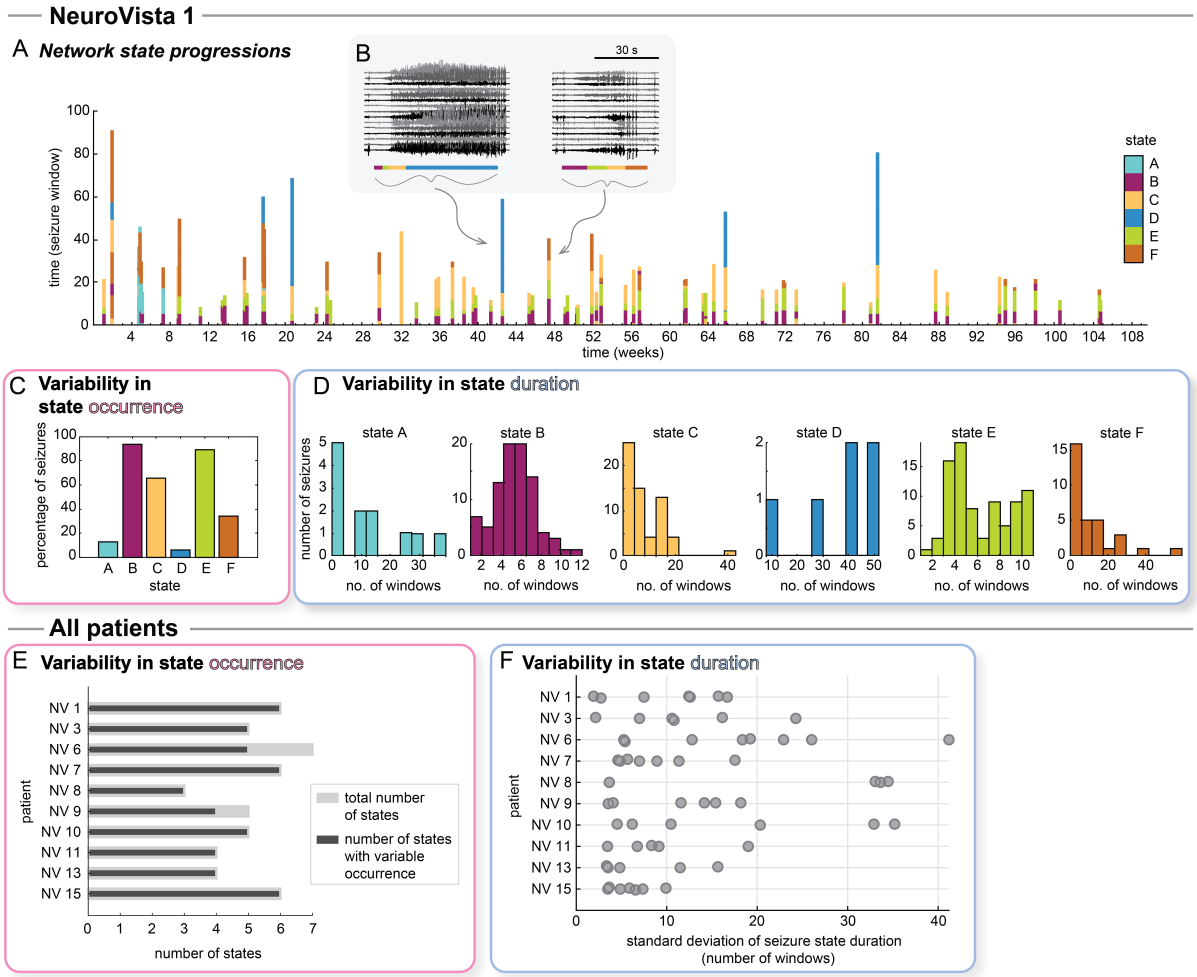

Figure 6.1: **Variability in seizure network state progressions.** Example patient, NeuroVista 1: A) Seizure network state progressions. Each seizure is represented by a vertical bar; its horizontal location indicates the time of seizure occurrence and the bar's height indicates the seizure's duration in terms of the number of time windows (approximately equal to number of seconds). The colours of the bar indicate the network state of the seizure during each time window. B) The iEEG traces of two example seizures, with the seizures' network state progressions shown underneath. C) Percentage of NeuroVisa 1's seizures that contained each seizure network state. D) Histograms of state durations of each of NeuroVista 1's network states. Seizures that do not contain a given state are excluded from the corresponding histogram (i.e., a state duration of zero is not included in the distribution). Across all patients: E) The number of seizure states in each patient (light gray bars), with the number of those states that had variable occurrence (i.e., did not occur in all of the patient's seizures) overlaid with dark gray bars. F) Variability in seizure state duration for each patient's states. Each dot corresponds to a single seizure state. NV = NeuroVista.

time windows a seizure spent in the given state) (Fig. 6.1D). For example, in NeuroVista 1, state D only occurred in about 6.4 % of seizures, and, in those seizures, its duration could range from 8 to 52 windows.

Across patients, we also observed variability in seizure network state progressions (Fig. 6.1E,F), although the patterns of variability varied from patient to patient (see Zenodo Data File 10.5281/zenodo.5910238 for seizure state progressions of all patients). Importantly, seizure states are not comparable across patients due to patient-specific iEEG implantations; however, for convenience, we allow states in different patients to share the same letter label and colour. All patients had variability in state occurrence for the majority of their seizure states (Fig. 6.1E). Indeed, only NeuroVista 6 and NeuroVista 9 had any states that occurred in all of their recorded seizures (two states and one state, respectively). Further, all seizure states had variable state duration, indicating that state durations were not a fixed feature across seizures (Fig. 6.1F). Thus, both state occurrence and state duration captured information about within-patient seizure variability in our cohort, and we therefore investigated patterns of variability in both features. State occurrence is a binary feature that described whether or not the state occurred in a seizure. Meanwhile, state duration is a continuous measure that, in seizures containing the state, captured how long the state persisted.

6.2.2 Seizure states vary over the duration of chronic iEEG recordings

We first asked if within-patient seizure state occurrence and duration varied over the timescale of each patient's chronic iEEG recording. Specifically, we explored whether seizures that occurred early in the recording had different features from those seizures that occurred later.

We first investigated relationships between the amount of time elapsed since the iEEG implantation and seizure state occurrence. In each patient and for each of the patient's seizure states, we divided the patient's seizures into two groups: seizures containing the state and seizures that did not contain the state. We then compared the recording times (i.e., time since implantation) of these two groups. Fig. 6.2A shows the relationship of an example state with recording time. In this patient, NeuroVista 13, the prevalence

of state C gradually increased over the recording period. While few seizures contained state C during the first few months of the recording, state C was present in the majority of seizures by the end of the recording. The temporal separability of seizures with and without state C can be characterised by the area under the curve (AUC) of distinguishing state occurrence based on seizure time, with an AUC below 0.5 indicating that the state preferentially occurred in earlier seizures and an AUC above 0.5 revealing that the state tended to occur in later seizures. Here, state C has an AUC of 0.71, which was significant after false discovery rate (FDR) correction for multiple comparisons (Wilcoxon rank-sum test, $p = 2.1 \times 10^{-12}$).

Across our cohort, eight out of the ten patients had at least one seizure state whose occurrence was significantly associated with the time since implantation (Fig. 6.2B). Ten states had AUCs greater than 0.5, while eight states had AUCs less than 0.5, revealing that a state's occurrence can increase or decrease across a recording. Most states occurred across the majority of the recording (Supplementary 6.5.3), indicating that the observed temporal associations were not driven by transient states that only occurred during the initial part of the recording. Instead, the pool of possible states remained relatively consistent across a recording, but the likelihood of observing each state changed over time, as in Fig. 6.2A.

In the previous section, we observed that seizures with the same states can additionally differ due to variable amounts of time spent in those states (Fig. 6.1D,F). We therefore next investigated whether seizure state duration was associated with time since implantation. For each state, we only analysed state duration in seizures containing the state, thus ensuring that our results were not driven by seizures without the state (i.e., where state duration would be zero). In NeuroVista 15, state D demonstrates how seizure state duration can vary over the length of the recording (Fig. 6.3A). Here, state D's duration was significantly higher in earlier seizures, as demonstrated by a significant Spearman's correlation ρ of -0.59 after FDR correction for multiple comparisons ($p = 2.4 \times 10^{-6}$).

Almost all patients had at least one state whose duration was either significantly positively (eight states) or negatively (eight states) correlated with time since implantation (Fig. 6.3B-C). These findings demonstrate that, across a patient's chronic iEEG recording, it

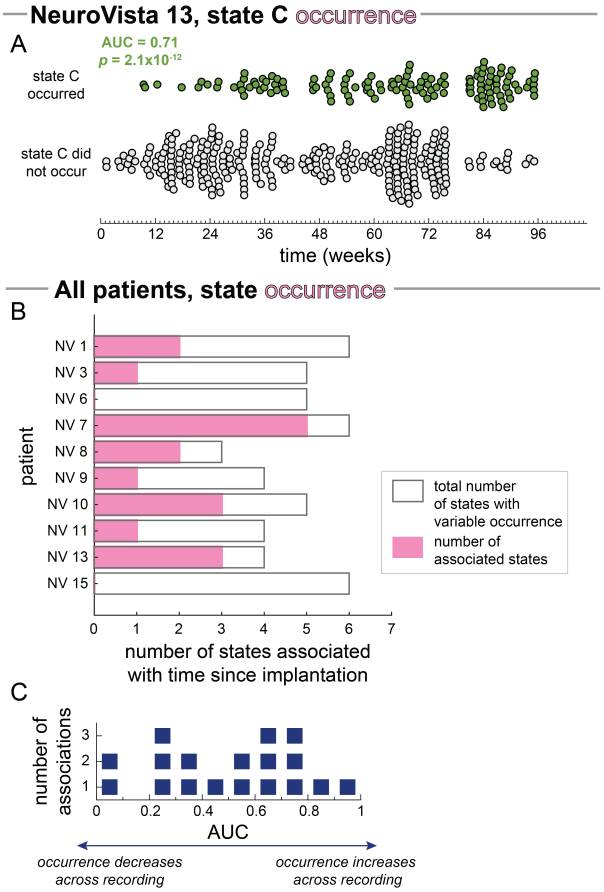

Figure 6.2: **Relationship between seizure state occurrence and time since implantation.** A) Time since implantation of seizures with (green circles) and without (grey circles) state C in example patient NeuroVista 13. Points are spread vertically to prevent overlap between seizures with similar times. Time since implantation separated seizures with and without state C with an AUC of 0.71. B-C) All patients: B) Number of seizure states whose occurrence was significantly associated with time since implantation in each patient (pink bars). Grey outlines provide a reference for the maximum possible number of associated states (i.e., the number of states with variable occurrence, equivalent to the dark grey bars in Fig. 6.1E). C) Dot plot of the AUCs for each significant state in (B). Each blue marker corresponds to one state.

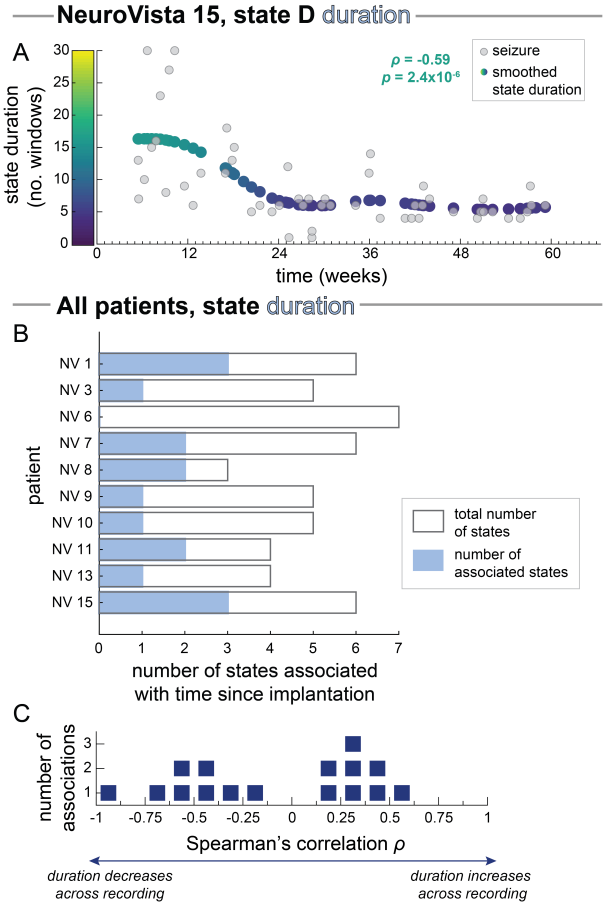

Figure 6.3: **Relationship between seizure state duration and time since implantation.** A) In example patient NeuroVista 15, state D duration versus the time since implantation, with a smoothed trend line (Gaussian window of 24 weeks) shown with the coloured points. Each grey point corresponds to a seizure that contained state D. B-C) All patients: B) Number of seizure states whose duration was significantly associated with recording time in each patient (light blue bars). Grey outlines provide a reference for the maximum possible number of associated states (i.e., the total number of states in each patient, equivalent to the light grey bars in Fig. 6.1E). C) Dot plot of the Spearman's correlation between state duration and recording time for all significantly associated states. Each blue marker corresponds to one state.

is also possible for a seizure state's duration to increase or decrease.

6.2.3 Seizure states fluctuate over circadian and multidien cycles

The previous results revealed seizure evolutions can differ across the multiple months of chronic iEEG recordings. However, we also hypothesised that seizure evolutions, like seizure occurrence (Baud et al., 2018; Karoly et al., 2021; Leguia et al., 2021; Maturana et al., 2020), may vary over circadian and multidien cycles. Importantly, these cycles can be nonstationary, with the cycle period varying over time; thus, they must be extracted using a continuous biomarker such as interictal spike rate (Baud et al., 2018; Karoly et al., 2021; Leguia et al., 2021; Maturana et al., 2020). We therefore next explored whether seizure state occurrence and seizure state duration were also associated with interictal spike rate cycles.

As in previous work (Baud et al., 2018; Karoly et al., 2016, 2021; Leguia et al., 2021; Maturana et al., 2020), we observed high levels of variability in interictal spike rate across each patient's chronic iEEG recording (see Fig. 6.4A for interictal spike rate of an example patient, NeuroVista 1). We obtained each patient's interictal spike times from a previous study (Karoly et al., 2016) and used a data-driven approach, EMD (Colominas et al., 2014; Huang et al., 1998), to extract fluctuations in spike rate over different timescales (see Methods for details and Zenodo Data File 10.5281/zenodo.5910238 for spike rate decompositions of all patients). EMD decomposes a time series into a series of oscillatory components that can have variable amplitude and frequency over time, thus accommodating non-stationarity in the spike rate cycles. To ensure that the extracted cycles were robust, we limited our analysis to the most prominent cycles that had short periods relative to the length of the patient's recording (recording length at least four times the average period of the cycle; see Methods for details). Fig. 6.4B shows the extracted spike rate cycles of NeuroVista 1. The fastest analysed cycle (yellow) was circadian, with an average period of 0.96 days. The patient also had a multidien cycle (red) of approximately 10.61 days. These spike rate cycles captured different timescales of changes in spike rate in NeuroVista 1's recording. We observed such multiscale fluctuations in interictal spike rate in most patients (Fig. 6.4C and D, see Supplementary section 6.5.2 for selection of spike

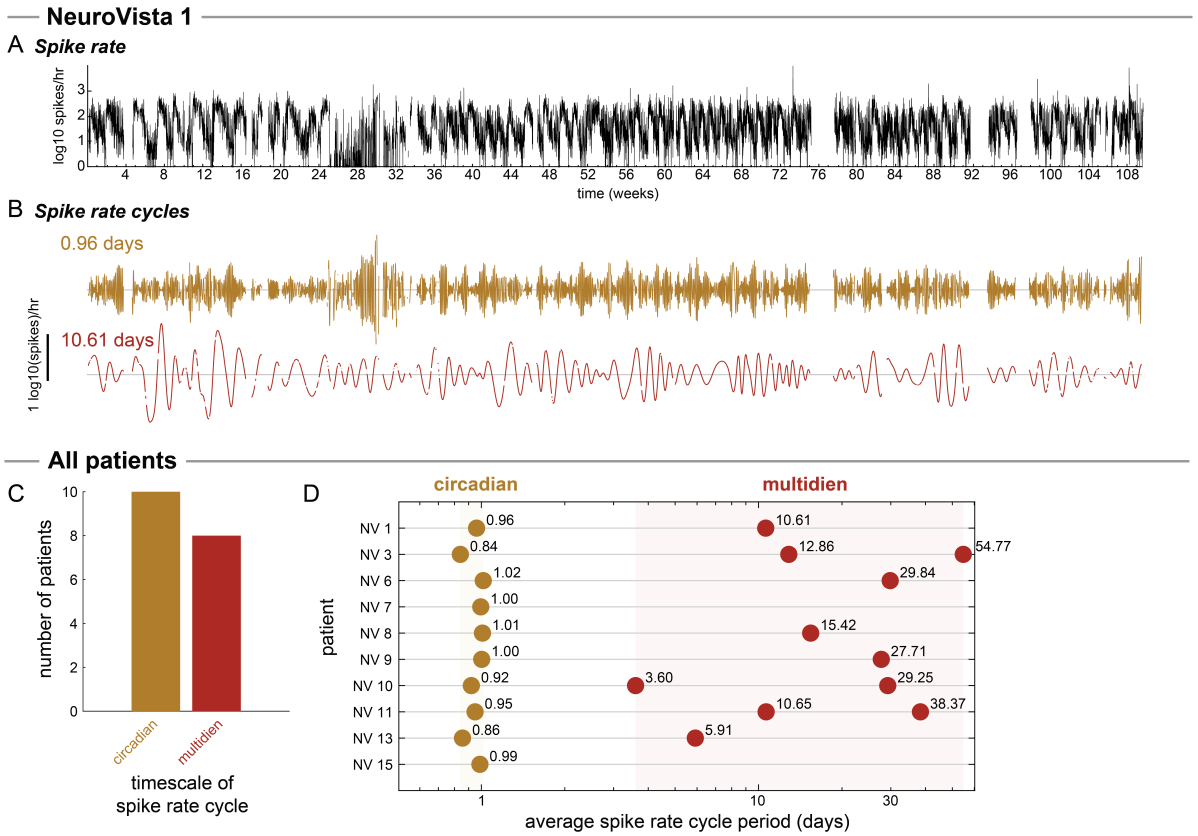

Figure 6.4: **Patient-specific cycles in interictal spike rate.** A-B) Example patient Neurovista 1: A) The interictal spike rate (number of spikes per hour versus recording time) during NeuroVista 1's recording. B) The two prominent cycles in NeuroVista 1's spike rate, extracted using empirical mode decomposition (EMD). NeuroVista 1 had a circadian (average period of 0.96 days) and multidien (average period of 10.61 days) cycle. C-D) All patients: C) Number of patients that had at least one spike rate cycle at each timescale. D) Spike rate cycles of each patient, coloured by their timescale.

rate timescales). All patients had prominent circadian cycles in spike rate (average period of 0.84-1.02 days), and eight of the ten patients also had at least one multidien cycle, with average periods ranging from 3.60 to 54.77 days. Together, these cycles characterised the prominent patient-specific, non-stationary changes in spike rate in each patient.

For each patient, we first asked whether seizure states preferentially occurred during certain phases of each spike rate cycle. Similar to previous work (Baud et al., 2018; Karoly et al., 2018a; Leguia et al., 2021), we defined phase preference as the phase locking value (PLV) of a state for a spike rate cycle. A PLV of 0 would indicate that the state had no phase preference, while a PLV of 1 would indicate that the state only occurred at one specific phase of the cycle. Importantly, seizures themselves usually have phase preferences

for circadian and multidien spike rate cycles (Baud et al., 2018; Karoly et al., 2021; Leguia et al., 2021; Maturana et al., 2020), which would by proxy lead to phase preferences for seizure states as well. Therefore, to control for seizure timing phase preferences, we used permutation tests to determine the significance of each state's PLV. In other words, we determined if the seizure state's phase preference was significantly stronger than the phase preference of the seizures.

Fig. 6.5A-C shows an example of a seizure state, state F, that preferentially occurred during certain phases of NeuroVista 1's multidien spike rate cycle. In this example, state F was most likely to occur during a specific part of the rising phase of the multidien cycle, with the proportion of seizures with this state tapering towards the cycle peak. Further, almost all seizures that occurred during the falling phase and the cycle trough lacked this state. As such, state F's PLV was significantly stronger than the overall seizure PLV (state PLV = 0.84, +0.21 relative to PLV of all seizures, $p = 0.0014$) after FDR correction for multiple comparisons.

Across our cohort, four patients (NeuroVista 1, 7, 10, and 13) had at least one seizure state whose occurrence was significantly associated with a spike rate cycle (Fig. 6.5D-E). Note that the same state could be associated with multiple different spike rate cycles (see Supplementary section 6.5.4 for associated states). Four of these patients had a state that had a phase preference in their circadian cycle, while two patients had one or more states associated with at least one multidien cycle (Fig. 6.5F). We also quantified the strength of the significant associations with spike rate cycles by computing their state PLV increases relative to the overall PLV of all seizures for the given cycles (Fig. 6.5G). The increase in PLV varied from 0.07-0.22, with most states showing a weak increase in PLV (median: 0.12). We interpret these associations as evidence that, in some patients, certain spike rate cycles reveal a modulation in seizure state occurrence over a specific timescale.

We then investigated if seizure state duration also varies over spike rate cycles. For seizures with a given state, we computed the non-parametric circular-linear correlation D (Mardia, 1976) between the seizures' state durations and the phases of the spike rate cycle at which the seizures occurred, using permutation tests to determine statistical significance (see Methods). The measure D is not signed, but instead ranges from 0 to 1,

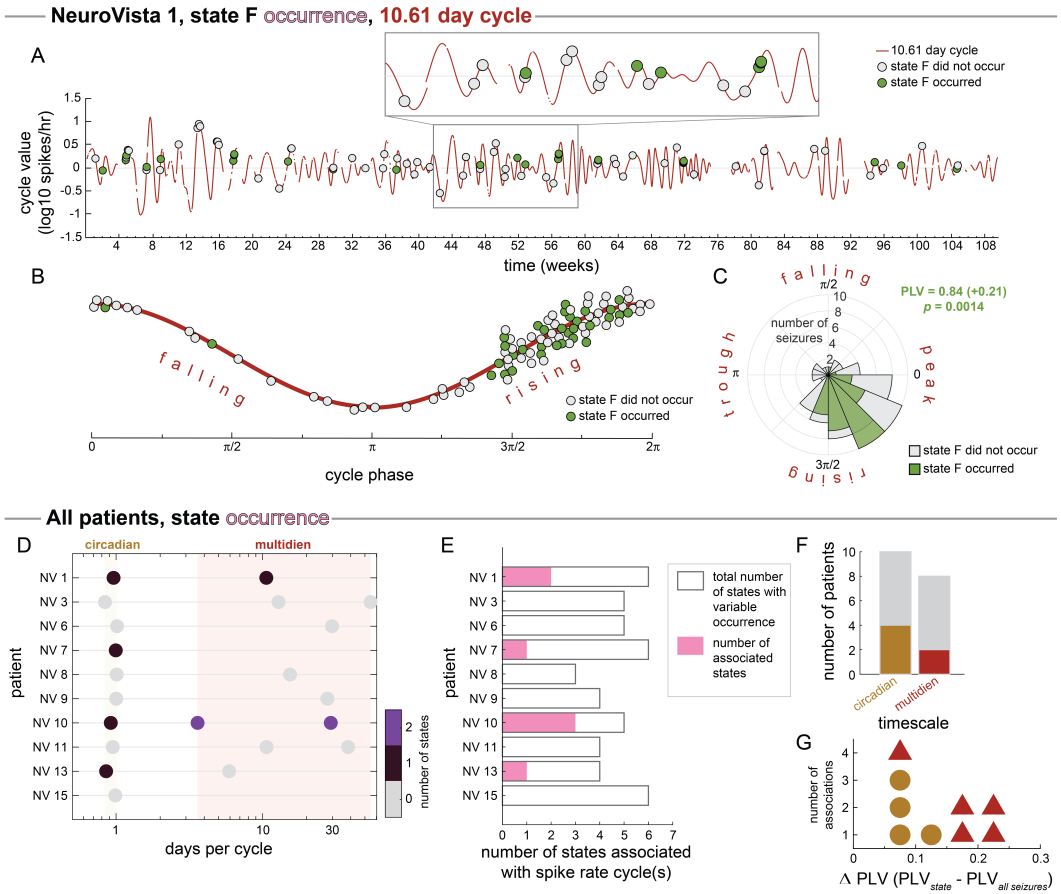

Figure 6.5: **Associations of seizure state occurrence with spike rate cycles.** NV = NeuroVista. A-C) Example patient NeuroVista 1, seizure state F, 10.61 day cycle: A) NeuroVista 1's 10.61 day spike rate cycle versus time, with seizures (circles) overlaid on the cycle at the times they occurred. Circles are coloured by whether the seizure did (green) or did not (grey) contain state F. The boxed inset expands a part of the cycle. B) Representation of the phase of the 10.61 day cycle, with the curve indicating whether spike rate is falling or rising at the given phase. As in A, circles correspond to seizures with (green) and without (grey) state F, and they are plotted on the curve at the phase that they occurred. Seizure circles are spread vertically to show all circles. C) Polar histogram of seizure phases with (green) and without (grey) state F. D-G) All patients: D) Coloured circles indicate the spike rate cycles. The circle colours correspond to the number of states significantly associated with the patient's spike rate cycle at that timescale. Grey circles indicate that the patient had a spike rate cycle at that timescale, but that that timescale was not associated with the occurrence of any seizure state. E) Total number of seizure states associated with spike rate cycles in each patient. Grey outlines provide a reference for the maximum possible number of associated states (i.e., the number of states with variable occurrence). F) For each timescale of spike rate cycle, the number of patients that had at least one state associated with that timescale. Grey bars indicate the number of patients with each timescale (equivalent to the coloured bars in Fig. 6.4C). G) Dot plot of the strength of the phase preferences of significantly associated states, measured as changes in PLV (state PLV minus the PLV of all seizures for the associated spike rate cycle). Marker shape and colour indicates the timescale category of the associated spike rate cycle.

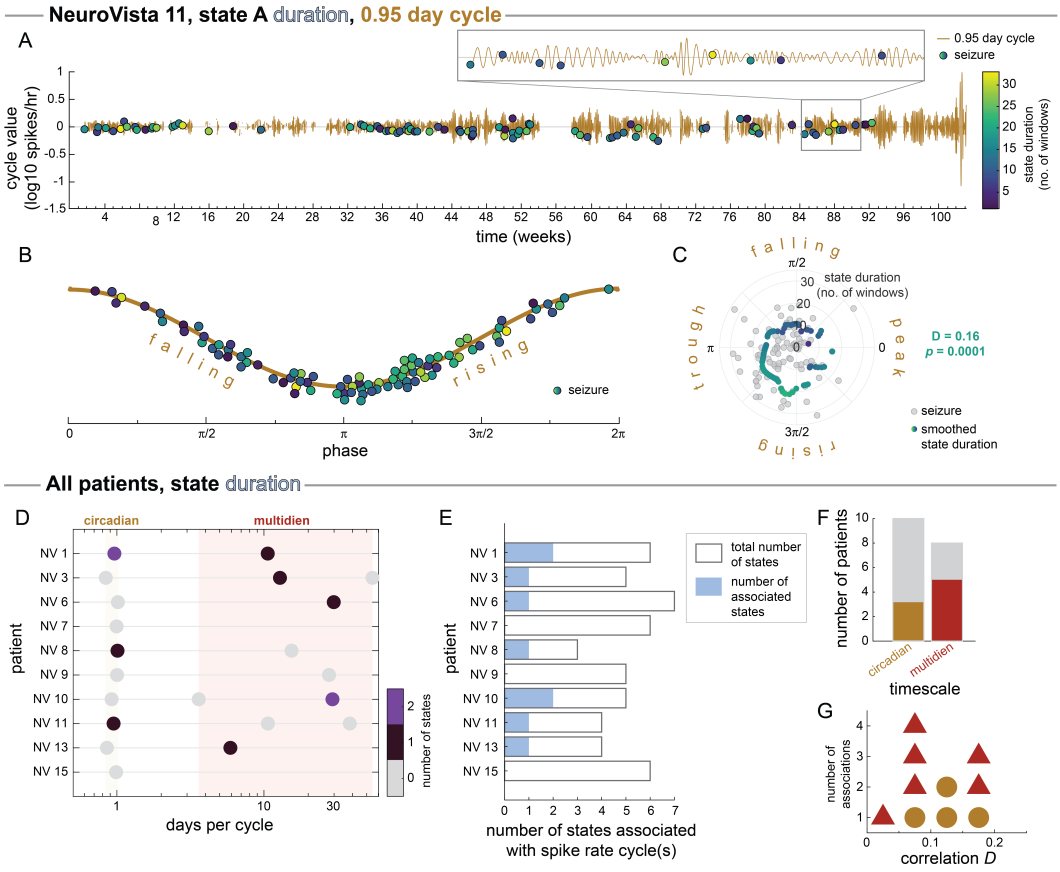

Figure 6.6: **Associations of seizure state duration with spike rate cycles.** NV = NeuroVista. A-C) NeuroVista 11, seizure state A, 0.95 day spike rate cycle: A) The circadian spike rate cycle versus time, with seizures containing state A (circles) overlaid on the cycle at the times they occurred. Circles are coloured by state duration. The boxed inset expands a part of the cycle. B) Representation of the phases of the spike rate cycle, with the curve indicating whether spike rate was falling or rising at the given phase. As in (A), circles correspond to seizures with state A, with their colour indicating state duration. Seizure circles are spread vertically to show all circles. C) Polar scatter plot of state duration versus seizure phase. Each grey point indicates a seizure, while coloured points correspond to the smoothed seizure duration ($\pi/4$ Gaussian window). D-G) All patients: D) Coloured circles indicate the spike rate cycles in each patient. The colours of the circles indicate the number of states significantly associated with the patient's spike rate cycle at that timescale. Grey circles indicate that the patient had a spike rate cycle at that timescale, but that that cycle was not associated with the duration of any seizure state. E) Total number of seizure states whose durations were associated with spike rate cycles in each patient. Grey outlines provide a reference for the maximum possible number of associated states. F) The number of patients that had at least one state duration associated with a spike rate cycle at each timescale. Grey bars indicate the number of patients with each timescale (equivalent to the coloured bars in Fig. 6.4C). G) Dot plot of the strengths of the significant associations between state duration and spike rate cycle phases, measured as non-parametric circular-linear correlation D. Marker shape and colour indicates the timescale category of the associated spike rate cycle.

with zero indicating no association. Seizures without the state (i.e., with a state duration of zero) were excluded from the analysis so that duration results were not driven by state occurrence phase preferences.

Fig. 6.6A-C shows an example seizure state duration association with a spike rate cycle in NeuroVista 11. In this patient, the circadian spike rate cycle was significantly associated with the duration of state A after FDR correction for multiple comparisons ($p = 0.0001$). State A's duration was markedly higher during the rising phase than during the falling phase of the spike rate cycle (Fig. 6.6B); the average duration starts increasing shortly before the trough of the cycle and peaks at approximately $3\pi/2$ in the rising phase before decreasing again (Fig. 6.6C).

In our cohort, seven of the ten patients had one or more states whose duration was significantly associated with spike rate cycle phases (Fig. 6.6D-E). Of these patients, three had circadian associations and five had multidien associations (Fig. 6.6F). The strength of the correlations between seizure state duration and spike rate cycle phases varied from 0.04 to 0.18, with most state durations showing a weak correlation with spike rate phase (median: 0.09) (Fig. 6.6G). As with seizure state occurrence, we interpret these associations as evidence that seizure state duration can be modulated over the timescales of spike rate cycles.

6.2.4 A seizure state's occurrence and duration are usually independently modulated

In Chapter 5, we found that seizure duration could vary independently of changes to the seizure's network evolution (Schroeder et al., 2021). This finding suggests that features that characterise seizure network evolutions (e.g., seizure state occurrence) and features that influence seizure duration (e.g., seizure state duration) could have different modulators. We therefore next asked whether a given state's occurrence and duration were associated with the same timescale, which we would interpret as co-modulation by the same time-varying process.

In each patient, we first determined whether each seizure state's 1) occurrence, but not duration, was associated with the time since implantation, 2) duration, but not occurrence,

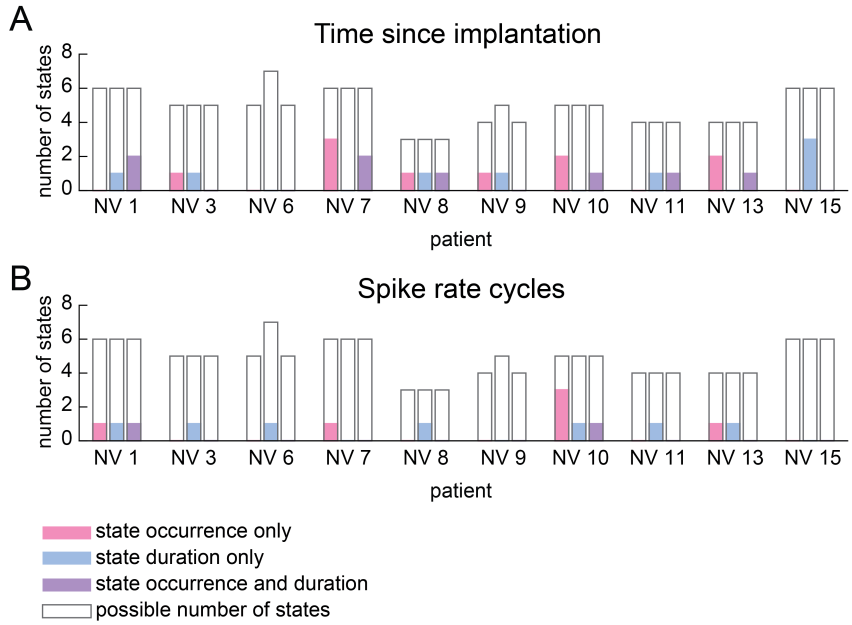

Figure 6.7: **Independent and coinciding timescales of seizure state occurrence and seizure state duration modulation** A) From left to right, the three bars for each patient show the number of the patient's seizure states in which 1) only state occurrence was associated with time since implantation (pink), 2) only state duration was associated with time since implantation (blue), and 3) occurrence and duration were both associated with time since implantation (purple). Grey outlines provide a reference for the maximum number of states that can belong to each category (i.e., for occurrence counts, the number of seizures states with variable occurrence, and for duration counts, the total number of seizures states) in each patient. B) Same measures as (A) for state associations with spike rate cycles.

was associated with the time since implantation, or 3) occurrence and duration were both associated with time since implantation. Fig. 6.7A shows the number of states in each patient that belonged to each category. Although the occurrence and duration of the same state were sometimes both associated with time since implantation, it was more common for only one to be associated with time since implantation.

We repeated this analysis for seizure state associations with spike rate cycles (Fig. 6.7B). Note that for this analysis, a state could belong to multiple categories if its occurrence and duration were associated with different spike rate cycles; for example, a state's occurrence could be associated with a circadian cycle (category 1) while its duration could vary over a multidien cycle (category 2). For spike rate timescales, a state's occurrence and duration were only associated with the same timescale in one state in NeuroVista 1 and one state in

NeuroVista 10 (Supplementary section 6.5.4). Thus, joint modulation of state occurrence and duration was rare over circadian and multidien timescales.

Finally, we investigated two relationships to support our overall results. First, we analysed whether certain parts of a seizure (e.g., onset states) were more likely to change over a given timescale. For both state occurrence and state duration, we did not find any significant associations between where the state occurred in the seizure evolution and whether the state was associated with the time since implantation or a spike rate cycle (Supplementary 6.5.5). These results indicate that multiple places in seizure evolutions are susceptible to temporal modulation. Second, we investigated whether total seizure duration was associated with the time since implantation and spike rate cycles (Supplementary section 6.5.4). We observed associations with total state durations in only a few patients, indicating that overall changes in seizure duration were not the sole driver of our observed state duration associations. In other words, seizure state durations did not simply vary due to changes in overall seizure duration, where increasing seizure duration proportionally increased the component states' durations. This observation supports our earlier hypothesis that specific parts of seizure evolutions may be more prone to temporal "elasticity" (Schroeder et al., 2021; Wenzel et al., 2017), or variable duration.

6.3 Discussion

We analysed variability in seizure network state progressions in chronic iEEG recordings, providing novel insight into the patterns and mechanisms of seizure variability. We found that in most patients, seizure states depended on when the seizure occurred in the recording, with some states becoming more or less prevalent and/or increasing or decreasing in duration as the recording progressed. Additionally, several patients had one or more states associated with circadian and/or multidien cycles in interictal spike rate. These associations suggest that seizure features are modulated over multiple timescales, including circadian and multidien timescales that can be revealed by interictal biomarkers.

We first found that seizure evolutions often depended on the amount of time that had elapsed since the start of the patient's recording. Variability over such long timescales (multiple months to years) may reflect non-cyclical changes in seizure evolutions due

to factors such as postimplantation effects (Chen et al., 2021; Ung et al., 2017, 2016), medication changes (Napolitano and Orriols, 2013), and slow changes in the epileptic network due to plasticity (Hsu et al., 2008). Alternatively, more gradual changes in seizure evolutions could be produced by cycles that have long periods relative to the length of these recordings. Analysing longer recordings could determine if persistent seizure variability reflects longer cycles, such as circannual cycles, in brain dynamics (Karoly et al., 2021; Leguia et al., 2021; Rao et al., 2020). Notably, transiently observed states at the beginning of recordings were uncommon in our cohort, suggesting that implantation effects rarely cause atypical seizure states. Thus, shorter presurgical recordings patients with epilepsy, which typically last for a few days to a few weeks, likely contain a patient's usual seizure states, although state duration and relative state prevalence may change over time. However, some additional seizure states may emerge after weeks of recording (King-Stephens et al., 2015; Ung et al., 2016) that would therefore be missed during presurgical monitoring. Thus, uncovering the longer timescales of variability in seizure evolutions has important implications for interpreting shorter iEEG recordings. Our findings add to the existing literature on variability in brain dynamics across chronic iEEG recordings (Chen et al., 2021; Ung et al., 2017, 2016) by revealing that multiple features of seizure evolutions also vary across these longer timescales.

We also observed associations between seizure state variability and circadian and/or multidien cycles in several patients. It is well-known that there are also cycles in seizure occurrence over the same timescales (Gowers, 1885; Karoly et al., 2016, 2018a; Langdon-Down and Brain, 1929; Loddenkemper et al., 2011; Navis and Harden, 2016; Patry, 1931), and recent research has revealed that fluctuations in interictal spike rate can serve as a biomarker for these seizure cycles (Baud et al., 2018; Karoly et al., 2021; Leguia et al., 2021; Maturana et al., 2020). We previously found evidence of circadian rhythms in seizure evolutions from short-term epilepsy monitoring unit recordings, and the same results suggested that slower trends also modulated seizure dynamics (Schroeder et al., 2020). Here, we confirmed that seizure features follow both circadian and multidien cycles that are linked to fluctuations in interictal spike rate. Compared to our previous study (Schroeder et al., 2020), our work here uses much longer recordings and larger numbers of seizures per

patient, providing new evidence for time-varying modulation of seizure evolutions. Moreover, we identified specific aspects of seizure evolutions, network state occurrence and network state duration, that changed over time in most patients. Our results reveal that there is a multiscale temporal structure to seizure variability, and certain characteristics of seizures are susceptible to modulation over a specific timescale.

We found that both seizure state occurrence and seizure state duration changed over the duration of the recording and were associated with spike rate cycles. Interestingly, a given seizure state's occurrence and duration were usually not associated with the same timescale, indicating that these features were modulated separately in most patients. This finding suggests that the absence of a state is not a special case of duration modulation where the state's duration is zero; in other words, the factors that decrease the duration of a state are likely different from the factors that prevent the state from occurring altogether. Additionally, it implies that a state's duration is not necessarily linked to its likelihood of occurring in the seizure; for example, a state does not necessarily persist for longer in seizures where the state is more likely to occur. Instead, separate factors appear to control whether a seizure state occurs and the duration of the state when it occurs. Additionally, we observed that the durations of specific seizure states had associations with time since implantation and spike rate cycles. This observation supports our earlier hypothesis that specific parts of seizure evolutions may be more prone to temporal "elasticity" (Schroeder et al., 2021; Wenzel et al., 2017), or variable duration from seizure to seizure (see Chapter 5). We suggest that future research on seizure duration variability likewise decompose seizure evolutions into states or other subparts to uncover factors that impact overall seizure duration.

Our work builds on past research that provided evidence for seizure variability over specific timescales. For example, it is well-established that in some patients, clinical seizure features such as secondary generalisation are associated with sleep/wake state or day/night cycles (Bazil, 2018; Bazil and Walczak, 1997; Janz, 1962; Loddenkemper et al., 2011; Sinha et al., 2006). Past analysis of chronic iEEG in canines also discovered shifts in seizure onset patterns as the recording progressed, likely due to postimplantation variability in brain dynamics (Ung et al., 2016). Additionally, variability in seizure onset and

spread have been linked to preictal and interictal changes in network features (Khambhati et al., 2016), band power (Naftulin et al., 2018), the location of high frequency oscillations (Gliske et al., 2018), and patterns of cortical excitability (Badawy et al., 2009). The same interictal features (network dynamics (Kuhnert et al., 2010; Mitsis et al., 2020), band power (Panagiotopoulou et al., 2022), high frequency activity (Chen et al., 2021), and cortical excitability (Meisel et al., 2015)) have all been shown to vary over circadian and/or multidien cycles. We now show that seizure evolutions also change over the timescales that influence interictal brain dynamics, suggesting that these fluctuations share common mechanisms.

Uncovering the mechanisms that influence these interictal and ictal dynamics is an active area of research. There is growing evidence that both seizures and spike rate are influenced by multiple factors operating over different timescales (Karoly et al., 2021; Rao et al., 2020). We therefore decomposed interictal spike rate into circadian and multidien cycles, rather than focusing on the relationship between overall spike rate and seizure features. Although past studies had sought a relationship between preictal spike rate and seizure likelihood, they found conflicting results, with decreases, increases, or no discernible patterns in preictal spike rate (Karoly et al., 2016). These inconsistent findings were clarified by later studies that decomposed spike rate and demonstrated timescale-specific phase relationships between spike rate cycles and seizure likelihood (Baud et al., 2018; Leguia et al., 2021; Maturana et al., 2020). For example, the circadian spike rate phase that is associated with seizure occurrence varies across patients (Baud et al., 2018; Leguia et al., 2021). These different circadian seizure "chronotypes" (Loddenkemper et al., 2011) likely arise due to patient-specific circadian influences on seizure likelihood, while spike rate is primarily driven by the sleep/wake cycle (Karoly et al., 2021; Rao et al., 2020). However, across patients, seizures are most likely to occur during the rising phase of the multidien spike rate cycle (Baud et al., 2018; Leguia et al., 2021). Notably, these phase preferences mean that seizures do not preferentially occur when spike rate is maximal. Additionally, seizures are less likely to occur during the falling phase of the multidien spike rate cycle, even though spike rate during the rising and falling phases will be comparable. These observations suggest that seizure occurrence depends on the hidden factors that influence

spike rate, leading to phase preferences that vary from factor to factor. Likewise, we observed specific phase relationships between seizure states and spike rate cycles (e.g., Fig. 6.6B,C, state duration highest during the rising phase). Our work therefore supports the hypothesis that the underlying factors that influence seizure likelihood and spike rate also shape seizure features.

To uncover the timescales of seizure variability, we first needed to extract quantifiable seizure features. In the previous chapters, we described seizure network evolutions as continuous pathways through network space and quantified the overall similarity of pairs of seizure pathways (Schroeder et al., 2021, 2020). These pathways provide the most accurate description of seizure evolutions by capturing all within-seizure variations in functional networks as well as all differences between seizure pathways. However, these seizure pathways and pathway dissimilarities do not readily translate to specific seizure features. In this chapter, we therefore clustered seizure functional networks into a small number of network states, transforming seizure pathways into seizure state progressions. We could then characterise each seizure by its states and state duration. Other papers have used similar approaches to describe seizure network evolutions (Burns et al., 2014; Khambhati et al., 2015) and found comparable numbers of within-patient seizure states. It is unclear if there is an optimal approach for clustering seizure networks, with past approaches including k-means clustering (Burns et al., 2014), non-negative matrix factorisation (NMF) (Frusque et al., 2020; Gliske et al., 2018; Khambhati et al., 2017), and community detection (Khambhati et al., 2015). It is likely that many approaches will yield similar seizure network states, especially for data with clear clusters. For our data, we found that NMF was well-suited for finding a representative network state for each seizure time window. Unlike hard clustering solutions such as k-means, the underlying NMF decomposition allowed for softer transitions between the seizure network states, a useful feature for decomposing seizure pathways without clear state transitions. Further, our use of stability NMF (Wu et al., 2016) ensured that we found robust seizure states and suggested that our seizure states were accurate representations of the seizure network data. Future work could explore if other clustering approaches yield similar results or apply more complex approaches, such as hidden Markov models (Vidaurre et al., 2017),

to define seizure states.

Our analysis of circadian and multidien timescales of seizure variability also relied on identifying cycles in interictal spike rate. To extract spike rate cycles, we used a data-driven method, EMD, that is adept at decomposing non-stationary time series (Huang et al., 1998). We additionally used a variation of EMD that improves the separation of cycles with different frequencies (Colominas et al., 2014). Unlike approaches such as wavelet decomposition, EMD does not require hypotheses about cycle frequencies or shapes, as cycles are instead found using the intrinsic fluctuations of the time series. EMD is also unique due to its iterative approach for extracting different timescales of cycles. This method provides advantages for spike rate analysis: the spike rate time series does not need to be smoothed to extract slower cycles, and cycles with drastically different periods can be uncovered without scanning a wide range of possible cycle frequencies. Although EMD differs from previous approaches for finding spike rate cycles (Baud et al., 2018; Karoly et al., 2018a; Leguia et al., 2021; Maturana et al., 2020), we found similar results, including circadian spike rate cycles in every patient as well as many multidien cycles. To increase confidence in our extracted spike rate cycles, we also limited our analysis to robust cycles that persisted across the recording and had relatively high contributions to the overall spike rate. Future work could use a less conservative approach and determine if less salient cycles also influence seizure evolutions.

Although we found many associations between seizure timing, spike rate cycles, and seizure states, we were unable to explain the full spectrum of seizure variability in our patients. Other approaches could yield more comprehensive and stronger explanations of seizure features. First, our analysis focused on spike rate phase due to its association with seizure occurrence (Baud et al., 2018; Karoly et al., 2021; Leguia et al., 2021). However, we also observed that the amplitude of spike rate cycles often varied over time, potentially reflecting variability in the strength of these cycles. Such changes in cycle strength could potentially impact seizure features. Second, as with seizure occurrence (Baud et al., 2018; Leguia et al., 2021; Maturana et al., 2020), different cycles likely interact to produce the observed seizure variability. A predictive model incorporating multiple timescales may be more informative than a single spike rate cycle (Panagiotopoulou et al., 2022). Third, we

limited our analysis to each patient's overall spike rate. Spatial patterns of spike rate also vary over time (Chen et al., 2021), and other interictal events such as high frequency activity have different temporal profiles than spike rate (Chen et al., 2021). Spatiotemporal variability in interictal dynamics may be linked to seizure variability (Gliske et al., 2018) and spatial patterns of different interictal events could also be incorporated in multivariate models of seizure features. Finally, although a given state's occurrence and duration were often independently modulated, there were likely interactions across seizure states. For example, the duration of an early state may depend on whether a seizure progresses to a subsequent possible state (Kaufmann et al., 2020). Approaches such as canonical correlation analysis (Zhuang et al., 2020) could uncover combinations of seizure features that are associated with combinations of interictal features.

Another open question is whether changes in the rate of seizure occurrence are linked to changes in seizure features. In particular, it would be interesting to explore if the timescales associated with seizure occurrence, which we did not examine in our study, also influence seizure evolutions. Many factors could impact both whether seizures occur and seizure features themselves. For example, increasing cortical excitability likely increases seizure likelihood (Meisel et al., 2015) and also facilitates seizure spread (Badawy et al., 2009; Enatsu et al., 2012). However, we also hypothesise that more subtle changes in spatial patterns of cortical excitability could change seizure evolutions without impacting seizure likelihood. Additional research is needed to determine if there are independent modulators and comodulators of seizure occurrence and seizure evolutions.

Understanding patient-specific seizure variability could provide new clinical strategies for managing seizures in patients with focal epilepsy. First, cycles in seizure features could be added to seizure forecasting algorithms (Baud and Rao, 2018; Freestone et al., 2017; Stirling et al., 2021), allowing them to forecast not only when a seizure will occur, but also how the seizure will manifest. Thus, seizure forecasting could help anticipate when seizures will be more severe and dangerous to the patient. Additionally, clinicians could modify a patient's antiepileptic medication based on both seizures likelihood (Baud and Rao, 2018; Ramgopal et al., 2013; Stirling et al., 2021) and seizure severity (Cramer and French, 2001). Both interictal and seizure variability may also have implications for treatment efficacy;

for example, in a mouse model of temporal lobe epilepsy, optogenetic stimulation only impacted seizures that arose from specific brain states (Ewell et al., 2015). Novel, seizure-specific treatments could therefore be designed to fluctuate over the same timescales as the patient's seizures, thus delivering time-adaptive treatments that account for the patient's seizure variability. Finally, uncovering the time-varying mechanisms that underlie seizure variability and severity could provide new targets for manipulating seizures and lessening their impact on patients. As a preliminary step towards clinical applications, future work could connect seizure network states to clinically relevant features, such as onset locations and secondary generalisation. Additionally, repeating our analysis in a larger cohort could determine if certain characteristic temporal patterns of variability, analogous to chronotypes in seizure occurrence (Langdon-Down and Brain, 1929; Leguia et al., 2021; Loddenkemper et al., 2011; Rao et al., 2020), exist across patients.

In summary, we have shown that features of seizure evolutions vary over multiple timescales within individual patients with focal epilepsy. Like interictal dynamics, seizures can change over the months and years of chronic iEEG recordings as well as over faster timescales, such as circadian and multidien cycles. As with cycles in seizure occurrence, cycles in seizure features can be extracted using interictal spike rate as a biomarker. Future work could explore whether fluctuations in other interictal features, such as spatial patterns of spikes and high frequency activity, explain additional seizure features. Ultimately, uncovering the timescales of within-patient seizure variability could lead to new time-adaptive approaches for controlling seizures.

6.4 Methods

6.4.1 Computing seizure functional network evolutions

We analysed long-term iEEG recordings from 10 NeuroVista patients (Cook et al., 2013). The iEEG of type 1 seizures (seizures with clinical manifestations and corresponding iEEG changes) and type 2 seizures (seizures with iEEG changes comparable to type 1 seizures, but without confirmed clinical manifestations) were extracted and preprocessed as described in Chapter 2. For each patient, the time-varying (sliding window) functional connectivity was then computed for each seizure and decomposed using NMF (see Chapter 2). Thus, each patient's seizure network evolutions were described by a small number of patient-specific NMF basis vectors and time-varying coefficients. Each NMF basis vector described functional network interactions between channels, while the NMF coefficients specified the contributions of each basis vector to the functional network of each seizure time window.

6.4.2 Computing progressions of seizure network states

We observed that most seizure time windows had a single NMF basis vector with a high coefficient. As such, a time window's dominant basis vector (i.e., the basis vector with the highest coefficient) provided a simplified description of the time window's functional connectivity. Therefore, a seizure's time-varying functional connectivity could be simplified as a sequence, or progression, of network states, where the network state of each time window was the dominant NMF basis vector (see Chapter 2). We used this approach to describe each patient's seizures as progressions of network states.

6.4.3 Preprocessing of interictal spike rate

Detection of interictal epileptiform spikes for this dataset was previously performed and validated (Karoly et al., 2016). The time-varying spike rate for each patient's recording was summarised as the number of spikes, across all channels, in non-overlapping one hour windows. Due to communication dropouts or failures to regularly store the iEEG data, each one hour segment could have missing segments that affected the spike rate count. To

normalise for these dropouts, we normalised each hourly spike rate count by the proportion of captured iEEG data:

$$S_{t,norm} = \frac{S_{t,rec}}{1-D_t}$$

where $S_{t,rec}$ is the recorded spike rate (spikes/hr) of hour t, $S_{t,norm}$ is the spike rate (spikes/hr) of hour t after normalising for dropouts, and D_t is the proportion of dropout time, or missing data, of hour t. Recording hours with $D \geq 0.75$ were considered missing data.

After normalising for dropouts, each hourly spike rate $SR_{t,norm}$ was log transformed:

$$S_t = log10(S_{t,norm} + 1)$$

yielding the final spike rate, S_t, of each hour t.

Beginning with the shortest missing segments of spike rate data, we then iteratively imputed missing segments of spike rate data using a method similar to Baud et al. (2018). For each missing segment, we first selected the spike rate data segments directly preceding and following the missing segment that were the same length as the missing segment. If this data was available (i.e., did not contain missing values or exceed the endpoints of the recording), we used the surrounding segments to generate spike rate data for the missing segment. This data was generated by linearly interpolating between the means of the surrounding segments and then adding Gaussian noise with a mean of zero and standard deviation of the surrounding segments. Any resulting interpolated data with a spike rate of less than zero was changed to zero. The length of each interpolated segment was also recorded so that interpolated data was only included in the analysis when its length was much shorter ($\leq 20\%$) than the period of the analysed spike rate cycle (extracted in Methods section 6.4.4). Any remaining missing time points were temporarily set to the mean spike rate value prior to EMD and then returned to missing values after the decomposition and Hilbert transform (Methods section 6.4.4).

6.4.4 Extracting interictal spike rate cycles using EMD

EMD (Huang et al., 1998) was then used to extract intrinsic mode functions (referred to as "spike rate cycles" in the Results) from each spike rate time series. Briefly, EMD decomposes any given signal into a set of signals known as IMFs that exactly reconstruct the original signal when summed together along with a with a residual signal. A key property of each IMF is that it must have approximately the same number (up to +/- 1) of extrema as zero-crossings, which ensures that there are no riding waves in the extracted IMFs, as well as a local mean (i.e., the mean of the maximal and minimal envelopes of the IMF) of zero. This property also ensures a well-defined Hilbert transform, allowing us to extract the phase of the signal fluctuations.

We used a variation of EMD known as complete ensemble empirical mode decomposition with adaptive noise (CEEMDAN) (Colominas et al., 2014) that helps ensure that each IMF contains oscillations with a similar timescale (i.e., the mode's period does not dramatically vary over time) by adding noise to the time series prior to the decomposition. The standard deviation of the added noise was scanned from 0.0025 to 0.125 in steps of 0.0025, and the decomposition at each noise level was performed with 100 noise realisations, a maximum of 1000 sifting iterations to extract each mode, and the signal-to-noise ratio increasing for every stage of the decomposition. This initial step yielded 50 versions, one for each noise level, of the EMD decomposition for each patient's spike rate time series.

For each decomposition, we used the Hilbert transform to determine the time-varying frequency, phase, and amplitude of each extracted spike rate IMF (Huang et al., 1998). We initially estimated the average period of each IMF using the median frequency of only the original (i.e., non-interpolated) spike rate data, excluding the first and last ten days of the recording due to possible instability in the frequency estimate at the time series boundaries. For each IMF, segments with interpolated spike rate were removed if their duration exceeded 20% of the IMF's period. To define each IMF's timescale, we then recomputed the average period of each IMF as above, now using all of the IMF's non-missing data. The average amplitude of each IMF was computed using the same process.

For each noise level, we then computed the pairwise index of orthogonality, O (Huang et al., 1998), between all pairs of time series from the EMD decomposition (i.e., the IMFs and the residue signal):

$$O_{i,j} = \frac{1}{T} \sum_{t=1}^{T} \frac{C_i(t)C_j(t)}{C_i(t)^2 + C_j(t)^2}$$

where $C_i(t)$ is the ith extracted time series, $C_j(t)$ is the jth extracted time series, t is the time point in each time series, and T is the total number of time points with spike rate data in both time series. The normalisation by $\frac{1}{T}$ allowed us to compare O across pairs of time series that had different amounts of missing data due to the spike rate interpolation step. O is close to zero when the two time series are locally orthogonal (i.e., do not contain oscillations at similar frequencies during the same time interval). Thus, to minimise overlap in the frequencies of different spike rate cycles, we found the maximum absolute value of the pairwise O for each decomposition and then selected the noise level that minimised this value. This decomposition was used for all downstream analysis.

The median amplitudes versus median periods of the IMFs of the selected decompositions are shown in Supplementary section 6.5.2. To focus our analysis on the primary, robust contributors to spike rate cycles, we limited our analysis to IMFs with locally prominent amplitudes that had median periods that were less than a quarter of the duration of the patient's recording (see Supplementary section 6.5.2). Across patients, we observed a clear distinction between cycles with median periods of approximately one day (0.83 to 1.03 days) and cycles with longer periods (3.93 to 54.77 days). We labeled these timescales as circadian and multidien cycles, respectively.

6.4.5 Comparing seizure state occurrence and seizure time since implantation

State occurrence was a binary seizure feature defined as whether a given state occurred in the seizure. Most seizure states did not occur in all of a patient's seizures; thus, these states had variable occurrence. For each patient and each of their seizure states with variable occurrence, a Wilcoxon rank sum test was used to compare the seizure times (i.e., the number of days after the recording's start that the seizure occurred) of seizures with

and without the state. To quantify the temporal separation of seizures with and without the state, the AUC of the receiver operating characteristic curve for distinguishing state occurrence using seizure times was also computed. Note that AUCs are mathematically equivalent to Wilcoxon rank sum tests, and therefore have the same statistical significance as the Wilcoxon rank sum test statistic.

6.4.6 Comparing seizure state duration and seizure time since implantation

State duration was a continuous measure that quantified the number of time windows a seizure spent in a given state. To compare state durations to seizure times in the recording, we computed the Spearman correlation between non-zero state durations and the number of days since the start of the recording that the corresponding seizures occurred. Spearman's correlation was also computed between the total duration of all seizures and seizure times in the recording.

6.4.7 Comparing seizure state occurrence and spike rate cycles

For each patient and each of their seizure states with variable occurrence, we determined the IMF ("spike rate cycle") phase preferences of seizures with the seizure state. For each IMF, a seizure's phase was defined as the phase of the IMF during the hour in which the seizure occurred. Seizures were excluded from the analysis if this spike rate data was missing. To quantify the phase preference of a given state and IMF, we first computed the mean resultant vector from the IMF phases of all seizures with the state:

$$Re^{-i\psi} = \frac{1}{S}\sum_{s=1}^{S} e^{-i\phi_s}$$

Here, S is the number of seizures containing the state, s is a seizure with the state, ϕ_s is the IMF phase of seizure s, R is the modulus of the mean resultant vector, and ψ is the angle of the mean resultant vector. As in previous work (Baud et al., 2018; Leguia et al., 2021), we refer to R as the PLV of seizures containing the state. The PLV varies from 0 to 1 and is higher when the IMF phases are similar across all seizures with the state.

To control for any seizure phase preference, we used permutation tests to determine the significance of the observed PLV for each seizure state and IMF. For a state that occurred S times, we randomly selected S seizures from all analysed seizures and recomputed the PLV. Repeating this process for 10,000 different permutations yielded a null distribution of PLV if the seizure state had no additional phase preference within that spike rate cycle. The p-value of the association was defined as the percentage of times a permuted PLV was greater than or equal to the observed PLV.

6.4.8 Comparing seizure state duration and spike rate cycles

For each patient, we also compared the durations of each seizure state to the patient's spike rate cycles. For each IMF ("spike rate cycle") and state, we found the IMF phases of seizures that contained the state and computed the rank linear-circular correlation D between the state's duration and the IMF phases (Mardia, 1976). D varies from 0 to 1, with higher values indicating a stronger association between the state's duration and the IMF phases. We determined the significance of these associations by permuting these state durations 10,000 times and computing a null distribution of correlations. The p-value of the observed correlation was the percentage of times a permuted correlation was greater than or equal to the observed correlation. We used the same approach to compare each patient's overall seizure durations to each IMF.

6.4.9 Comparing spike rate and seizure states

We also compared the overall spike rate, prior to EMD, to state occurrences and state durations using Wilcoxon rank sum tests and Spearman's correlation, respectively (Supplementary section 6.5.4).

6.4.10 Correction for multiple comparisons

Benjamini-Hochberg FDR correction for multiple comparisons (Benjamini and Hochberg, 1995), with $\alpha = 0.05$, was performed for all tests across all patients that compared seizure features (state occurrence, state duration, and seizure duration) to temporal features

(seizure time in the recording, spike rate cycles, and overall spike rate). Uncorrected p-values are reported in the text.

6.4.11 Code and data availability

All analysis was performed in MATLAB version R2018b. CEEMDAN of interictal spike rate was performed using the MATLAB package CEEMDAN (`https://github.com/macolominas/CEEMDAN`) (Colominas et al., 2014; Torres et al., 2011). The remaining analysis was performed using custom MATLAB scripts. The NeuroVista seizure iEEG data used in this study is available from `www.epilepsyecosystem.org`. The processed data (NMF W and H matrices) of all patients is available on Zenodo (DOI 10.5281/zenodo.5503590).

6.5 Supplementary

6.5.1 Seizure network states of an example patient, NeuroVista 1

Fig. 6.8 shows the six network states of NeuroVista 1. Each state was derived using stability NMF (Wu et al., 2016) (see Chapter 2).

6.5.2 EMD of interictal spike rate and selection of spike rate cycles for comparison with seizure features

For each patient, we used EMD to extract cycles in interictal spike rate. Fig. 6.9 shows the median amplitudes versus median periods (days/cycle) of each EMD cycle. We limited our analysis to cycles that had both (1) locally prominent amplitudes, which we defined as local maxima (found using MATLAB function *findpeaks*) in these plots with amplitudes greater than or equal to 20% of the patient's highest median IMF amplitude and (2) median periods of less than a quarter of the recording's duration. These cycles are shown using the coloured circles in each patient. The fastest IMF of each patient was not included in the analysis since these oscillations are often thought to contain temporally unstructured noise.

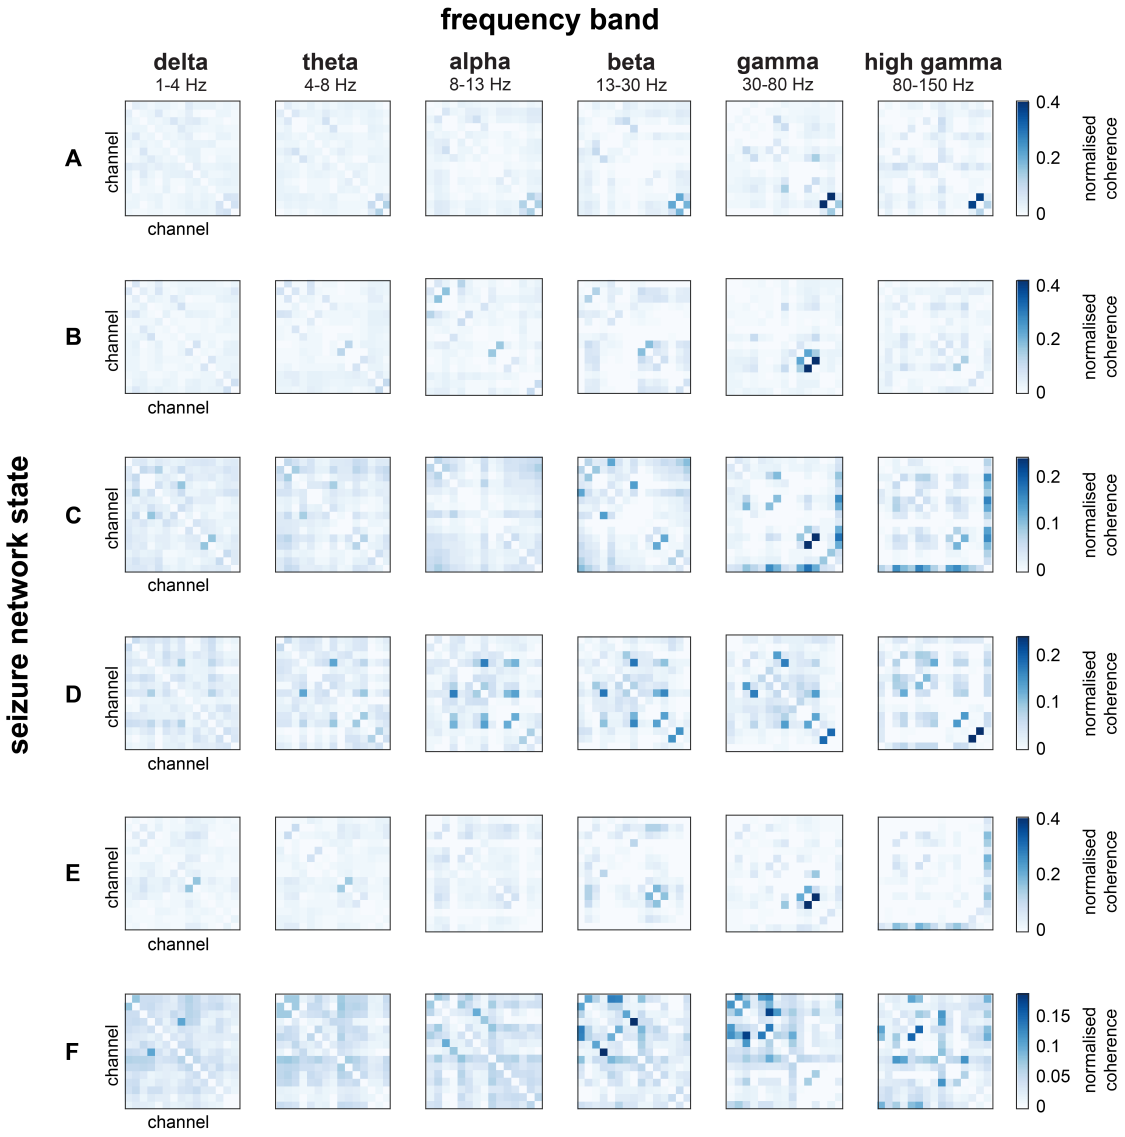

Figure 6.8: **Seizure network states of NeuroVista 1.** Rows correspond to seizure network states. Each network state describes the coherence between pairs of iEEG channels in six frequency bands (columns). Thus, each network state is composed of six functional connectivity matrices, one for each frequency band. Colormap limits are consistent within each state (i.e., across each row). Self-connections (diagonal matrix elements) are not shown.

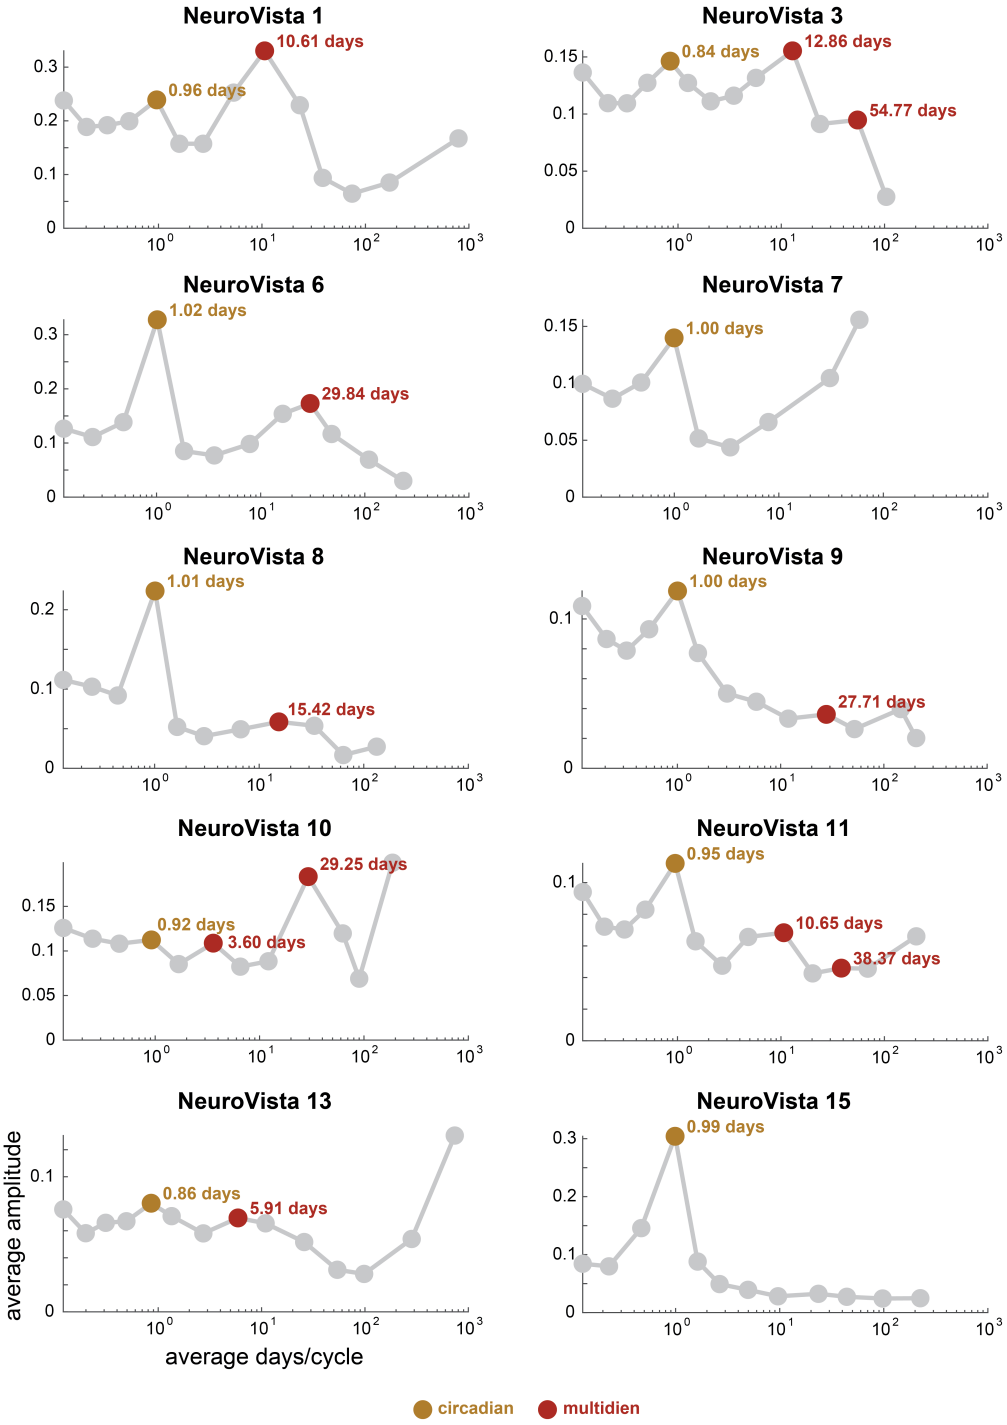

Figure 6.9: **Amplitudes and periods of the spike rate cycles extracted usng EMD.** Each plot shows the median amplitude and period of all of the extracted spike rate cycles (circles) of each patient. Coloured circles indicate cycles that were analysed in this study, with the colour indicating the timescale category of the cycle (circadian or multidien). Spike rate residues are not shown.

6.5.3 First and last occurrences of seizure states in patient recordings

Fig. 6.10 shows the first and last known occurrences of each seizure state in each patient. Each state visualisation is also coloured by whether the state's occurrence was significantly associated with the time since the start of the recording. As a reminder, states are patient-specific, even if they share the same letter label. In most cases, the time from first to last occurrence of a seizure state spans the majority of the patient's recording. Thus, most states are not limited to a specific section of the recording.

Note that due to missing and noisy data, some of the patient's seizures were not captured in the iEEG recordings. As such, some states may have occurred earlier or later in the recording period than displayed here. As such, this analysis can only conclusively say when seizure states *did* occur, and there is some uncertainty regarding when states *did not* occur.

6.5.4 Supplementary spike rate, seizure duration, and seizure state analyses and visualisations

Fig. 6.11 shows which seizure state and total seizure duration were significantly associated with seizure time since implantation, spike rate cycles, and spike rate in each patient. We emphasise that states are *not* comparable across patients, even if they share the same label and colour.

The first columns in Fig. 6.11A and 6.11B show which states were significantly associated with overall spike rate after FDR correction for multiple comparisons. While overall spike rate was often associated with these seizure features, it was usually associated with the same or fewer states than specific timescales (spike rate cycles and time since implantation) (Fig. 6.11). The only exceptions were the occurrence of state C in NeuroVista 1, the durations of states C and D in NeuroVista 6 and, the duration of state C in NeuroVista 10, which were only associated with overall spike rate. Cases where overall spike rate, but not individual spike rate fluctuations, are associated with states may be due to spike rate features that our analysis did not capture (e.g., cycle amplitude) or joint effects of modulations over different timescales that were not significant at the level of individual timescales.

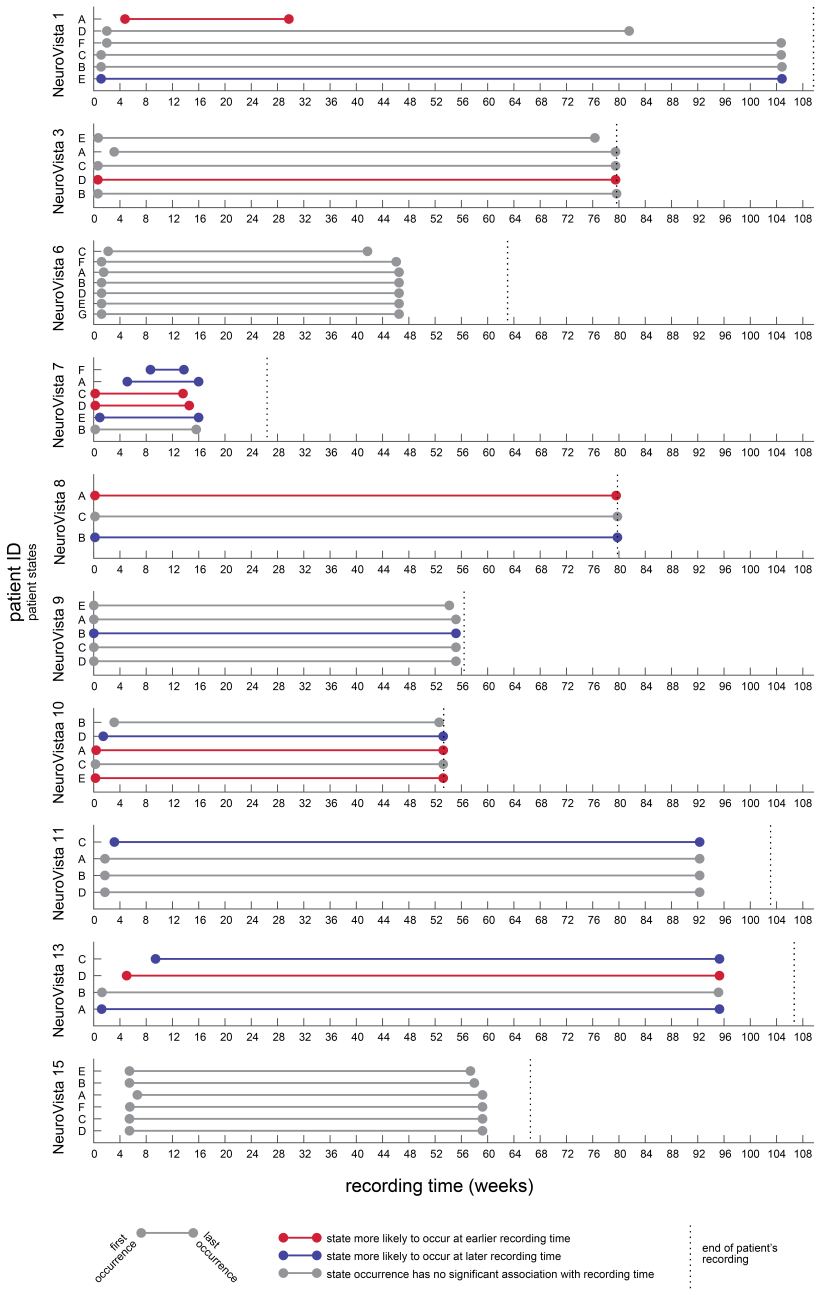

Figure 6.10: **First and last known occurrences of each seizure state in each patient.** In each patient, the times in the recording of the first and last occurrences of each seizure state (circles), with the time spanned by those occurrence marked with horizontal line. States are ordered from the shortest to longest amount of time that they spanned. State markers are coloured by whether the state's occurrence was significantly associated with the time since implantation.

Fig. 6.11B also indicates when a patient's total seizure duration was associated with overall spike rate, spike rate cycles, and/or recording time. Seizure duration was only associated with spike rate cycles in two patients and recording time in four patients; thus, state duration associations were more widespread than seizure duration associations in our cohort.

6.5.5 Locations of modulated states in seizure network state evolutions

We also investigated whether modulated states tended to occur during certain parts of seizure evolutions. In each patient, we first identified the order in which states occurred in each seizure. Only the first occurrence of a state was considered; for example, in the hypothetical seizure state progression

$$BBBAAACCACCCEEEEEE$$

state B would be first, A would be second, C would be third, and E would be fourth. This seizure does not provide any information about other states, such as state D. A state's *location* was then defined as the state's most common (i.e., mode) order of occurrence across all the patient's seizures that included the state. Fig. 6.12A shows the distributions of state locations in each patient. Across all patients, we then computed the mean state location for four categories of states (Fig. 6.12B, top row, left to right): (1) states whose occurrence was associated with seizure recording time, (2) states whose duration was associated with seizure recording time, (3) states whose occurrence was associated with at least one spike rate cycle, and (4) states whose duration was associated with at least one spike rate cycle.

We then performed permutation tests to determine whether significant states occurred earlier or later in the seizure evolutions than expected by chance (Fig. 6.12B, bottom row). For each scenario, we permuted state labels in each patient, selected the same number of "significant" states and their corresponding locations from each patient, and then recomputed the mean state location across all patients. Thus, this permutation test accounts for the distribution of state locations in each patient. For each test, the null distribution of mean state locations was computed using 10,000 permutations. The p-value

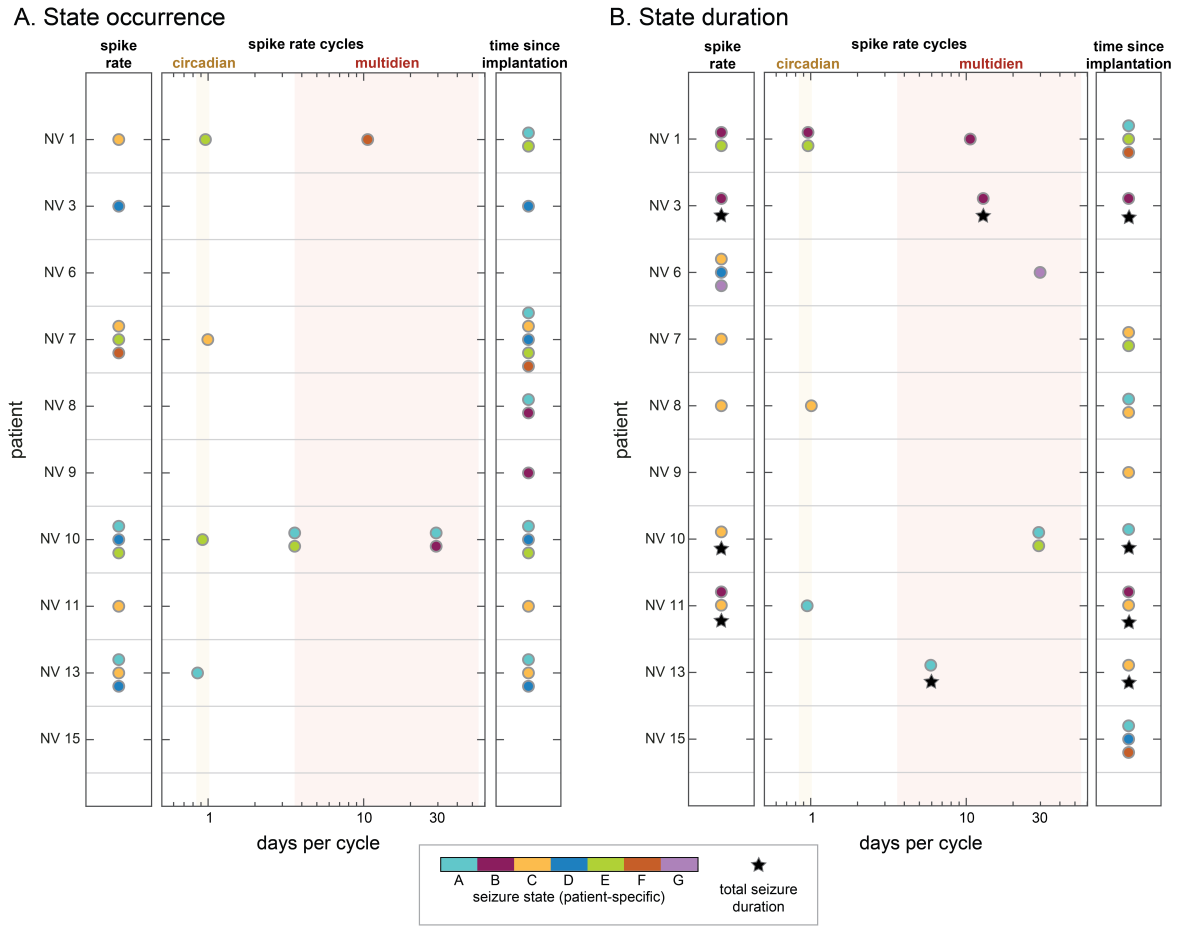

Figure 6.11: **Seizure states significantly associated with recording time, spike rate cycles, and spike rate.** Coloured circles indicate the seizure states with occurrences (A) or state duration (B) were significantly associated with overall spike rate (left columns), spike rate cycles (middle columns), and recording time (right columns). Rows for different patients are demarcated by horizontal grey lines. States are not comparable across patients, but are comparable within each patient. In (B), associations with total seizure duration are also marked with a black star.

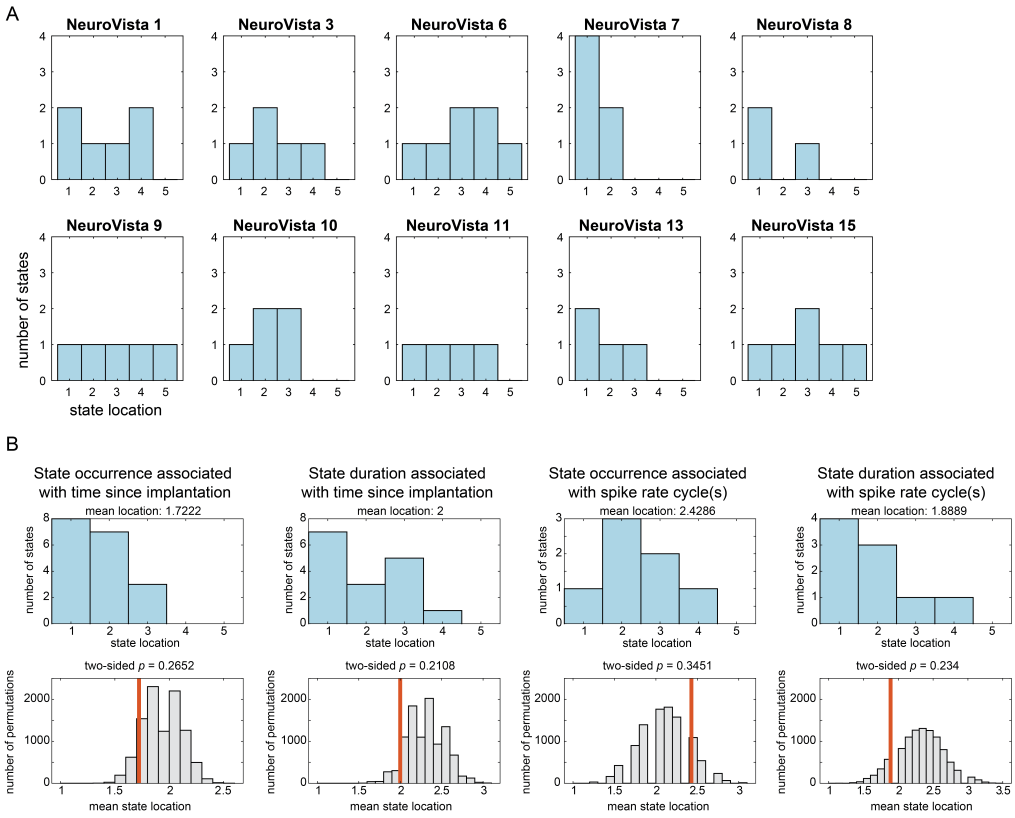

Figure 6.12: **Locations of significant seizure states in seizure network state evolutions.** A) Distributions of state locations in each patient. B) State locations of states significantly associated with seizure recording time and spike rate cycles. Top row: distributions of locations of significantly associated states. Bottom row: permutation test results for determining the significance of the mean state location in each scenario. The observed mean state location is marked with a red line in each distribution of permuted mean state locations.

of the observed mean location was defined as the percentage of permutations with a mean location as extreme (i.e., equidistant from the centre of the distribution) as the observed mean location. In all scenarios, the mean state location did not significantly differ from chance, suggesting that states that occurred earlier or later in seizure evolutions were not preferentially modulated.

6.5.6 Detecting signal dropouts

The NeuroVista data contains time periods of signal dropouts when the iEEG signal was not recorded. We used line length to identify iEEG segments with no signal (i.e., a flat time series with no voltage changes). We defined the line length L of a time series as

$$L = \frac{1}{T-1} \sum_{i=1}^{T-1} |x_{i+1} - x_i|$$

where x_i is the ith time point in a time series with T time points.

The time-varying line length of each seizure was computed for each iEEG channel in sliding windows (1/10s window, 1/20s overlap). Any time windows with 8 or more channels with line length ≤ 0.5, along with the preceding and following time windows, were considered missing data.

Chapter 6: Interictal Spike Rate Reveals Timescales of Modulation in Seizure States and Seizure State Durations

Chapter 7. Discussion: Implications of Within-patient Seizure Variability and Directions for Future Research

Contents

7.1 Contributions of This Thesis

Although epilepsy is a dynamic disorder in which symptoms fluctuate over time, it is unclear whether and how variability in its most salient feature - seizures - can arise within the same patient. Our understanding of within-patient seizure variability has largely been limited to select seizure features, such as seizure duration (Cook et al., 2016) and the extent of seizure spread (Karthick et al., 2018; Marciani and Gotman, 1986; Naftulin et al., 2018), that are known to change from seizure to seizure in some patients. However, fluctuations in seizure characteristics could have important implications for seizure control (Ewell et al., 2015; Martinet et al., 2015; Ryzi et al., 2015; Spencer et al., 1981) and seizure forecasting (Cook et al., 2016). As such, a comprehensive analysis of the prevalence and characteristics of such variability has been needed. In this thesis, I addressed this gap by providing new insight into within-patient seizure variability in patients with focal epilepsy. My analysis focused on comparing seizure functional network evolutions, which I conceptualised as seizure "pathways" through the space of possible brain network dynamics. Using this approach, I analysed seizure pathways in three different cohorts and addressed three open questions about within-patient seizure variability:

1. How prevalent is variability in seizure pathways within individual patients?

2. Does a seizure's pathway also determine it duration, or is there independent variability in these features?

3. How do seizure features change over different timescales?

In Chapter 3, I first addressed the need for an objective approach for comparing seizure pathways. My pairwise "seizure dissimilarity" measure built on previous methods for comparing seizures (Le Bouquin-Jeannès et al., 2002; Wendling et al., 1997, 1996, 1999; Wu and Gotman, 1998) by recognising seizures that share the same pathway, even if they progress at different rates. This temporal warping also easily allowed me to compare seizures with different durations without truncating the longer seizure. Using this measure, I compared seizure functional network evolutions in 31 epilepsy monitoring unit (EMU) patients with

approximately 2-16 days of intracranial EEG (iEEG) recording, revealing seizure variability in all patients. Interestingly, in most patients, this variability resulted from a spectrum of seizure pathways, rather than the coexistence of clear seizure types with different pathways. I also found that this variability did not simply arise from the presence of different International League Against Epilepsy (ILAE) clinical seizure types within the same patient. Overall, this chapter provides a new approach for comparing seizures and reveals a level of within-patient seizure variability that has not been previously acknowledged in computational epilepsy research.

After uncovering the prevalence of within-patient seizure variability, the next question that I began to address in Chapter 4 was how seizure pathways changed over time within each patient. In particular, it was unclear if seizures in each patient changed randomly over time or if there was some sort of temporal structure to when different pathways occurred. In the same cohort of EMU patients, I found that over the timescales of these recordings, seizures with more similar pathways tended to occur closer together in time. As such, within-patient changes in different seizure pathways are not random. Additionally, analysing temporal correlations between seizures over different timescales also revealed changes consistent with circadian fluctuations in seizure pathways. In almost all patients, the observed temporal patterns of seizure variability could be explained by a combination of circadian and/or slower time-varying factors that influenced seizure pathways. These findings suggest that modulatory processes shape within-patient seizure pathways over multiple timescales.

These first two research chapters focused on variability in seizure pathways. However, seizure dynamics are characterised not only by their pathways, but also by how long it takes to complete those pathways - i.e., their durations. In Chapter 5, I explored how variability in seizure pathways and seizure durations interact in the same patient, addressing open questions such as, "Do similar seizure pathways also have similar durations, or can there be temporal 'elasticity' (Wenzel et al., 2017) within the same pathway?" and, "Do populations of seizures with different durations correspond to distinct seizure pathways?" In this chapter, I analysed seizures in EMU patients as well as chronic iEEG recordings in both human and canine subjects. While most patients had a weak to moderate re-

lationship between seizure pathways and seizure durations, there was also independent variability in each feature: seizures with approximately the same pathway could have drastically different durations, and seizures with similar durations could traverse distinct pathways. Interestingly, in patients with duration populations, seizure duration was only a reliable proxy for seizure pathways in about half of patients. Overall, this analysis revealed that seizure pathways and durations are not tightly linked, suggesting that they can be independently modulated. Additionally, seizure pathways cannot be reliably inferred from seizure duration information alone, even in patients with distinct duration populations. Instead, both pathway and duration information is needed to fully characterise most seizures' spatiotemporal evolutions.

In my final research chapter, Chapter 6, I built on the insights from the previous chapters to explore how seizure features change over multiple timescales. For this analysis, I used chronic iEEG recordings from 10 patients with focal epilepsy. To quantify changes in seizure features, I characterised seizure pathways as progressions of a small number of network states, allowing me to separately analysis spatial features (whether each state occurred in a seizure) and temporal features (the duration of each seizure state) of seizure pathways. This approach therefore captures the different types of variability - pathway variability and duration variability - observed in Chapter 5. I found that each of these features were associated with fluctuations in interictal spike rate over circadian and multidien timescales, suggesting that seizure pathways are modulated by time-varying factors. Additionally, seizure states often changed across the months or years of a patient's recording. A given state's occurrence and duration were usually not associated with the same timescale, indicating that these features are separately modulated. These findings provide further evidence that time-varying factors shape seizure features within individual patients and, for the first time, uncovered longer timescales of changes in seizure pathways.

Thus, across these chapters, I have presented a new approach for comparing seizure pathways, demonstrated the prevalence of within-patient variability in seizure pathways, revealed that seizure pathways and durations can vary independently, and uncovered multiscale temporal fluctuations in seizure features. Together, this work provides a new perspective for analysing seizures within individual patients and suggests that seizures

themselves are modulated by time-varying factors. In the rest of this chapter, I will discuss areas of potential future work on seizure variability within individual patients with focal epilepsy, with a focus on exploring the types of variability in seizure pathways, uncovering underlying mechanisms that influence seizure features, and determining the clinical implications and applications of my work.

7.2 Mapping Within-patient Variability in Seizure Pathways

Early in my thesis, I proposed an approach for comparing seizure pathways that produces a single measure for the (dis)similarity of a pair of seizures. As such, this measure does not reveal exactly how seizure pathways differ within the same patient, and further work is needed to characterise the types of variability in seizure pathways. Here I present a conceptual model for how seizure pathways can differ and discuss potential approaches for exploring types of seizure variability.

7.2.1 A hypothesised model for producing variability in seizure pathways

In the previous chapters, I observed that two seizures could (1) traverse approximately the same pathway through network space, (2) share part of the same pathway, or (3) traverse completely distinct pathways, with no overlap in network space. However, the prevalence of these different scenarios across patients and the ways that they can arise are uncertain. For example, are some seizure pathways truncated versions of other seizure pathways, created by a subset of seizures terminating earlier along the full possible pathway (Karoly et al., 2018b; Wagner et al., 2015)? Can two seizures share the same initial pathway, but then diverge? Understanding how seizures can vary and the mechanisms that cause pathways to diverge could reveal opportunities for therapeutically controlling seizures.

To illustrate some of the potential ways that seizures could differ, Fig. 7.1A presents a conceptual model of how variability could arise in seizure pathways. For simplicity, I describe this model using the state progression simplification of seizure pathways (see Chapter 6), but the concepts also apply to continuous descriptions of seizure pathways. In this model, there is a framework that constrains possible seizure pathways in the patient.

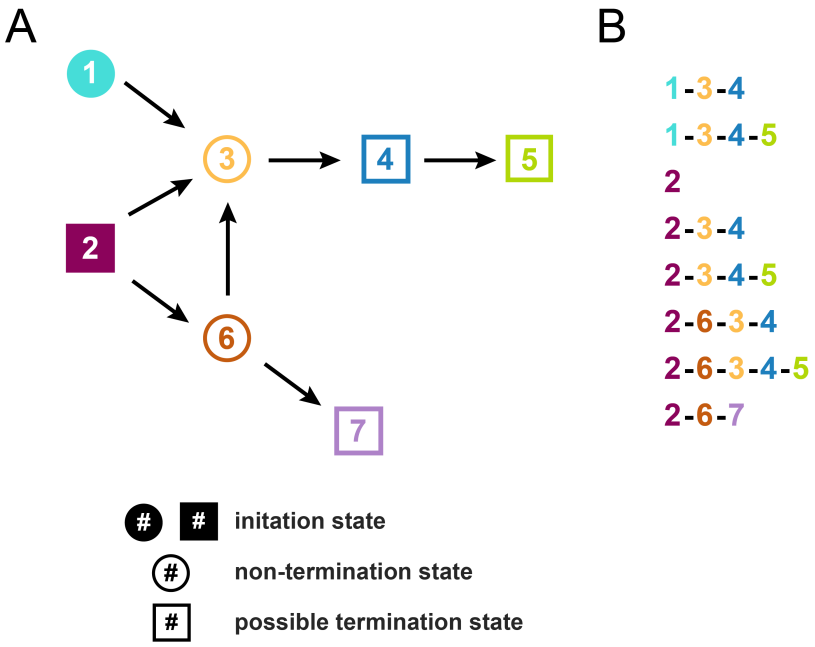

Figure 7.1: **Hypothesised model for generating variability in seizure pathways.** A) Diagram of possible seizure pathways in a hypothetical patient. Seizure pathways are described as progressions of network states. States that are filled in (states 1 and 2) are possible initiation states for a seizure pathway. Square states indicate points in the progression where the seizure may terminate. While some transitions are deterministic (e.g., state 3 always progresses to state 4), other states are "decisions points" at which variability is introduced into the seizure progression. Variability can be introduced by alternative onsets (e.g., onset states 1 and 2), different possible progressions (e.g., state 6 can progress to either state 3 or 7), and potential termination points (e.g., state 4 can terminate the seizure or progress to state 5). B) Possible seizure pathways arising from this model. The pathways demonstrate variability in state onset, state progression, and state termination. The probability of different onsets, transitions, and terminations would determine the relative prevalence of these different pathways.

However, the paths that a seizure can take are not fully deterministic; instead, "decision points" where a seizure pathway can evolve in different ways leads to variability among the patient's seizure pathways. For example, here seizures can start from different states, progress to different states from the same initial state, and terminate at different places along a possible progression. As such, this model provides a "map" of the possible seizure pathways within an individual patient (Fig. 7.1B).

This model contrasts with some past perspectives of seizures as deterministic events that, once started, follow a pre-determined course (Cook et al., 2016; Schindler et al., 2011).

However, it does not suggest that seizure pathways are random. Instead, within a patient, seizures are constrained to certain pathways, analogous to the way water follows existing stream and river beds. These shared dynamical elements across seizures explain why within-patient seizures can appear similar to each other, as has been previously observed (Burns et al., 2014; Karoly et al., 2018b; Kramer et al., 2010; Schevon et al., 2012; Truccolo et al., 2011; Wagner et al., 2015). Nonetheless, a single seizure pathway cannot summarise the full repertoire of the patient's seizure dynamics.

7.2.2 Dynamic maps of seizure pathways

The above model provides a static map of possible seizure pathways. However, as revealed in Chapters 4 and 6, within-patient seizure pathways change over time and are likely modulated over various timescales. As such, the maps of potential seizure pathways within individual patients are likely dynamic, with the probabilities of different decision points fluctuating over time and changing the prevalence of different seizure pathways. For example, a specific pathway may only be accessible during a certain phase of a circadian cycle.

Additionally, maps of seizure pathways could potentially change over longer timescales such as years and decades due to alterations in the underlying epileptic pathology. Earlier, I compared seizure pathways to streambeds and riverbeds that determine the "flow" of seizures. Such water channels are not static; they are also shaped by the flow of water, which in turn changes the water flow itself. It is unclear whether seizure activity or other abnormalities likewise change a patient's possible seizure pathways through mechanisms such as plasticity or structural alterations (Cole, 2000; Galovic et al., 2020; Huang et al., 1998). Alternatively, changes due to non-pathological processes such as ageing could also result in unidirectional shifts in a patient's map of potential seizure pathways. Further analysis in longer recordings, along with longitudinal measurements of hypothesised influential factors, would be needed to elucidate long-term changes in seizure pathways and the underlying mechanisms.

7.2.3 Approaches for mapping seizure pathways

While I anecdotally observed different patterns of variability in seizure pathways, further work is needed to test and refine the proposed model. One approach would be to create patient-specific models of seizure states and state transitions. A patient's seizure pathways could then be summarised as diagrams akin to the one in Fig. 7.1A with an accompanying transition matrix that describes probabilities of different onsets, transitions, and terminations. The probability of staying in each state, rather just transition probabilities between states, could also be included to capture state durations. Burns et al. (2014) used a similar approach to compare within-patient seizure state progressions, although they created individual transition matrices for each seizure, rather than summarising transitions across seizures. As Burns et al. (2014) also suggested, these seizure states could be used to create Markov Models of patient-specific seizure evolutions (Dadok et al., 2015).

Another approach to quantitatively mapping seizure variability would be to use variations of dynamic time warping (DTW) and other time series analysis methods to analyse the continuous time series that describe seizure pathways. For example, there are DTW variations that partially match time series (Giorgino, 2009) that could be used to find subsets of conserved dynamics across seizures pathways. The DTW time series warpings could also reveal parts of pathways that tend to have variable durations across seizures. To map these possible pathways, future studies could compute vector fields that summarise seizure trajectories through network space (Stitt et al., 2017). Importantly, small changes in brain state are known to change seizure onset and offset dynamics by qualitatively changing how the brain transitions between seizure and non-seizure periods (Saggio et al., 2020). Unlike the state description of seizure pathways, a vector field analysis could likewise reveal whether subtle early differences in seizure pathways alter the seizure's downstream evolution. For example, a small difference in seizure onset could potentially lead to an earlier seizure termination. Uncovering such cases could reveal when and how to intervene in seizures so that the seizure's travel route and destination have less impact on the patient.

Importantly, both of the above approaches could be computed on sliding window subsets

of seizure data to create a series of time-varying maps of seizure pathways. Depending on the timescale, these dynamic maps could capture cyclical or directed changes in seizure pathways. Dynamic maps of seizure variability may help reveal which seizure features and parts of seizure pathways are modulated over different timescales.

7.3 Hypothesised Mechanisms for Seizure Variability

One of the main unanswered questions raised by my work is what mechanisms shape seizure features and produce variability in seizure dynamics over time. Uncovering mechanisms of seizure variability will likely require a combination of testing and generating hypotheses using computational models, monitoring and probing the neural and broader biophysical states of patients with epilepsy, and manipulating neural dynamics in experimental models. Crucially, understanding how and why seizures naturally fluctuate over time could provide more detailed information for seizure forecasts as well as potential mechanisms for therapeutically mitigating seizure spread and severity. In this section, I will discuss some mechanisms that likely influence seizure evolutions within individual patients and approaches for testing these hypotheses in future work.

7.3.1 Spatiotemporal patterns of cortical excitability

Cortical excitability is often characterised by how strongly cortical neurons respond to a stimulus (Badawy et al., 2009, 2010, 2013; Freestone et al., 2011). More generally, cortical excitability depends on the balance of excitation and inhibition in the brain, which is thought to be pathologically altered in epilepsy (Badawy et al., 2012). Levels of cortical excitability are associated with seizure occurrence (Badawy et al., 2009; Meisel et al., 2015) and control (Badawy et al., 2010), and patterns of cortical excitability may influence the extent of seizure spread (Badawy et al., 2009; Enatsu et al., 2012). Uncovering associations between spatial patterns of cortical excitability and other seizure features may require monitoring this measure at finer spatial and temporal resolutions, either by continuously probing cortical excitability (Freestone et al., 2011) or developing localised measures of intrinsic excitability (Meisel et al., 2016, 2015) that do not require actively probing the brain's state.

Importantly, specific alterations in excitation and inhibition may be linked to changes in certain seizure features. For example, certain types of inhibition may control the rate of seizure propagation, and thus seizure duration (Trevelyan et al., 2007; Wenzel et al., 2017). Computational models can be used to explore the effects of different changes in neural excitation and inhibition (Jirsa et al., 2014; Wang et al., 2015, 2017). These mechanisms can also be tested pharmacologically in experimental models (Trevelyan et al., 2007; Wenzel et al., 2017) to determine the cellular and molecular bases of seizure variability.

7.3.2 Interictal and preictal spatiotemporal brain dynamics

Beyond direct measures of cortical excitability, a number of other interictal and preictal features may be linked to seizure variability. As discussed in the previous chapters, patterns of band power (Naftulin et al., 2018)), functional connectivity (Khambhati et al., 2016), and high frequency oscillations (Gliske et al., 2018) have been linked to changes in seizure features. One hypothesis is that preictal network dynamics may foreshadow functional networks that will be active during seizures (Khambhati et al., 2017). Alternatively, even if seizure and preictal network dynamics do not have similar structures, seizures with similar pathways could all arise from periods with similar preictal dynamics. However, not all preictal features necessarily impact seizure pathways; instead, only select features may be informative for anticipating seizure features. Indeed, the relevant preictal variability in brain dynamics could be subtle compared to other, non-pathological fluctuations (Mitsis et al., 2020). As such, the challenge may be uncovering which subsets of preictal features influence seizure evolutions. Approaches such as canonical correlation analysis (Zhuang et al., 2020) could be useful for extracting important preictal features and distinguishing modulators of different seizure features, such as pathway and duration variability.

Importantly, as in seizure forecasting, the timescales of fluctuations in brain dynamics may also be crucial for understanding changes in seizure features. For example, in Chapter 6, I found that phases of specific spike rate cycles were associated with more seizure features than overall changes in spike rate. In the field of seizure prediction and seizure forecasting, there has been a shift from trying to identify preictal signatures to instead

uncovering fluctuations that determine seizure risk (Stirling et al., 2021). Inconsistent links between seizure occurrence and interictal features (e.g., spike rate and signatures of impending critical transitions) have been clarified by extracting cycles in these features, rather than limiting the analysis to the preictal period (Baud et al., 2018; Karoly et al., 2021; Leguia et al., 2021; Maturana et al., 2020; Mitsis et al., 2020). Likewise, analysing seizure variability from the perspective of multiscale influential timescales, rather than preictal periods, could reveal more robust associations between interictal features and seizure variability (Panagiotopoulou et al., 2022). Specific modulatory timescales could arise if the interictal feature is a biomarker for other underlying mechanisms (Rao et al., 2020) or if seizure pathways are more sensitive to long-term *changes* in brain dynamics than to the brain state itself.

7.3.3 Variable spread on structural networks

Epilepsy is also associated with changes in whole-brain structural connectivity (Englot et al., 2016), which describes the pairwise connections between different brain regions. Further, alterations in structural connectivity are associated with across-patient differences in seizure onset location (Besson et al., 2014), seizure spread (Sinha et al., 2021a), and surgical outcome (Sinha et al., 2021b). However, it is less clear whether and how structural connectivity, which is relatively static over short timescales, impacts seizure variability within individual patients. One possibility is that structural connectivity constrains seizure spread, limiting possible seizure pathways. Faster, fluctuating spatiotemporal dynamics could then determine how seizures manifest given the structural limitations. Determining the relationships between structural connectivity, interictal/preictal brain activity, and seizures will likely require incorporating patient-specific structural networks into computational models of seizures (Proix et al., 2018), comparing structural connectivity and time-varying functional connectivity in interictal, preictal, and ictal periods (Shah et al., 2019), and comparing seizure propagation patterns and structural connectivity. Additionally, tracking longitudinal changes in structural connectivity alongside seizures could determine if slower changes in structural networks (Park et al., 2020) shift seizure evolutions over long timescales.

7.4 Clinical Implications of Seizure Variability

Seizure variability appears to be a common feature in patients with refractory focal epilepsy, regardless of the seizure onset area or patient pathology (Chapter 3). While the level and features variability do not appear to distinguish different pathologies, they could have implications for current treatments and the development of novel approaches for controlling seizures.

7.4.1 Antiepileptic medication

Antiepileptic drugs (AEDs) are one of the primary treatments for epilepsy and act on targets that influence neural excitation and inhibition (Macdonald and Kelly, 1995). While past studies have explored how AEDs relate to interictal dynamics and seizure frequency (Badawy et al., 2010; Meisel et al., 2016, 2015), little is known about whether and how AEDs change seizure features. Exploring whether and how seizures change in response to AEDs could be potentially provide an alternative measure of AED efficacy and could reveal if AEDs disproportionally impact certain seizure features or pathways. This specificity may arise from regional differences in the structure of neuronal circuits (Semyanov, 2003) and the distributions of AED targets such as GABA(A) receptors (Olsen, 2018; Trimmer and Rhodes, 2004). Such studies could pave the way for AED chronotherapy (Baud and Rao, 2018; Ramgopal et al., 2013) in which AED doses and/or types are adapted to fluctuations in seizure pathways. For example, a higher AED dose may be required to control more severe seizures. Ultimately, chronotherapy could both improve seizure control and reduce side effects by delivering the minimum effective dose needed to control seizures at a given time.

7.4.2 Surgical resection and surgical outcome

For patients with focal epilepsy whose seizures do not adequately respond to AEDs, surgical resection of the region responsible for generating seizures, the epileptogenic zone, is often the next line of treatment (Rosenow and Lüders, 2001). Using seizure functional networks to localise epileptogenic tissue and predict the outcome of surgical resection has

been an active area of research (Burns et al., 2014; Goodfellow et al., 2016; Kini et al., 2019). Since most of these studies analyse only a small number of seizures in each patient and/or average information across seizures, it is unclear how within-patient seizure variability in EMU data (Chapters 3 and 4) impacts these analyses. In particular, in cases where the patient has a single epileptogenic zone, do different seizure network evolutions contain redundant localisation information? If not, how can we tell which seizures will be most informative for presurgical planning? For example, it is unclear whether subclinical seizures are relevant for presurgical planning (Farooque and Duckrow, 2014). Some subclinical seizures and clinical seizures have similar evolutions (Chapter 3), and thus may contain similar localisation information. However, subclinical seizures with atypical spread, relative to clinical seizures, are associated with poor surgical outcomes (Farooque and Duckrow, 2014). As such, some subclinical seizure evolutions may contain additional information about the patient's pathology.

A particular concern for surgical planning is whether a patient has multifocal epilepsy, in which seizures are generated from multiple brain regions (Rosenow and Lüders, 2001; Ryzi et al., 2015; Spencer et al., 1981), or an area that could become epileptogenic (Rosenow and Lüders, 2001), leading to poor seizure control after surgery. I found that the level of seizure variability itself is not a predictor of surgical outcome (Chapter 3), and further work is needed to determine which types of seizure variability are signs of multifocal epilepsy. For example, it is unclear if seizure functional networks that give disparate localisations of the epileptogenic zone would multifocal epilepsy.

Another open question is what parts of seizure network evolutions are most informative for localisation. Interestingly, Burns et al. (2014) found that secondarily generalised seizures tended to have two periods that could help localise the putative epileptogenic zone: one shortly after seizure initiation, in which the epileptogenic zone is isolated from the rest of the functional network, and another during the later part of the seizure, in which the epileptogenic zone becomes well-connected to other areas. This second location could occur at different relative times across seizures, which is unsurprising given the observed temporal elasticity in seizure pathways (Chapter 5). As such, identifying informative *places* along pathways, rather than relative *times* in seizures, may be a more reliable

method for selecting seizure data for localisation.

7.4.3 Seizure predicting and seizure forecasting

Accurately anticipating seizures (seizure prediction) or periods of high seizure risk (seizure forecasting) could greatly mitigate the impact of seizures on patients' daily lives (Arthurs et al., 2010; Stirling et al., 2021). Incorporating information about within-patient seizure variability in seizure prediction and forecasting work could provide two main benefits. First, seizures with different features may have different preictal signatures (Cook et al., 2016). In these patients, algorithms that recognise multiple preictal signatures or find common elements of disparate preictal signatures may be more accurate. Additional research is needed to determine which types of seizure variability have implications for seizure prediction. For example, the seizure dissimilarity measure presented in Chapter 3 could be used to determine whether seizures missed by algorithms are distinct from the patient's other seizures.

Second, seizure prediction and forecasting could be extended to anticipate not only seizures, but also the characteristics of the seizure. In Chapter 6, I demonstrated that, like seizure occurrence, some seizure features appear to be modulated over specific timescales. Thus, these cycles could potentially be used to forecast seizure features as well. Since some seizures present higher risks of accidents (Lawn et al., 2004), this information could allow patients to further adapt their behaviour based on their seizure forecasts. Additionally, treatments such as AEDs and neurostimulation could be adapted based on the forecast to improve the control of specific seizure pathways.

7.4.4 Neurostimulation

Brain stimulation that prevents or quickly terminates seizures is emerging as an alternative treatment for patients with epilepsy (Jarosiewicz and Morrell, 2021; Sisterson et al., 2019; Wang et al., 2015). There is particular interest in closed-loop systems that detect seizure onset and rapidly use stimulation to prevent further seizure spread or, alternatively, use stimulation to steer the brain away from pro-ictal states. As for seizure forecasting, within-patient seizure variability has two main implications for this treatment. First, seizures may

respond differently to the same stimulation depending on their characteristics and their preictal state (Ewell et al., 2015). As such, complete seizure control may require different stimulation protocols that target different seizure pathways. Second, understanding the natural modulators of seizure pathways could suggest mechanisms for controlling seizure spread. For example, the stimulation could trigger an existing mechanism for earlier termination of a seizure pathway. Thus, a seizure-specific approach to neurostimulation could ultimately improve seizure control.

7.5 Conclusion

In focal epilepsy research, patient-specific models and analyses of epileptic processes have yielded many insights and opportunities to improve clinical treatments (Burns et al., 2014; Freestone et al., 2017; Goodfellow et al., 2016; Jiménez-Jiménez et al., 2015; Kini et al., 2019; Lagarde et al., 2019; Martinet et al., 2015; Proix et al., 2018; Sinha et al., 2016; Wang et al., 2019). More recently, there has been shift towards accounting for not only heterogeneity across patients, but also variability within individual patients due to fluctuations in brain dynamics over time (Baud et al., 2018; Chen et al., 2021; Gliske et al., 2018; Karoly et al., 2016, 2018a, 2021; Leguia et al., 2021; Panagiotopoulou et al., 2022; Proix et al., 2021; Saggio et al., 2020). My work here has added to this field by uncovering how seizures themselves vary over time within individual patients with drug-resistant focal epilepsy. Further research is needed to fully map patient-specific seizure pathways, determine the underlying mechanisms that shape seizure features, and utilise this knowledge to improve patient outcomes. These avenues provide exciting opportunities to further understand fluctuations in epileptic dynamics and design adaptive treatments that fully control seizures.

Bibliography

Afra, P., Jouny, C. C., and Bergey, G. K. (2015). Termination patterns of complex partial seizures: An intracranial EEG study. *Seizure*, 32:9–15.

Aggarwal, C. C., Hinneburg, A., and Keim, D. A. (2001). On the surprising behavior of distance metrics in high dimensional space. *Database Theory – ICDT 2001*, pages 420–434.

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

Ali, F., Rickards, H., and Cavanna, A. E. (2012). The assessment of consciousness during partial seizures. *Epilepsy and Behavior*, 23:98–102.

Arthurs, S., Zaveri, H. P., Frei, M. G., and Osorio, I. (2010). Patient and caregiver perspectives on seizure prediction. *Epilepsy and Behavior*, 19:474–477.

Badawy, R., Macdonell, R., Jackson, G., and Berkovic, S. (2009). The peri-ictal state: Cortical excitability changes within 24 h of a seizure. *Brain*, 132:1013–1021.

Badawy, R. A., Macdonell, R. A., Berkovic, S. F., Newton, M. R., and Jackson, G. D. (2010). Predicting seizure control: cortical excitability and antiepileptic medication. *Annals of Neurology*, 67(1):64–73.

Badawy, R. A. B., Freestone, D. R., Lai, A., and Cook, M. J. (2012). Epilepsy: Ever-changing states of cortical excitability. *Neuroscience*, 222:89–99.

Badawy, R. A. B., Jackson, G. D., Berkovic, S. F., and Macdonell, R. A. L. (2013). Cortical excitability and refractory epilepsy: a three-year longitudinal transcranial magnetic stimulation study. *International Journal of Neural Systems*, 23(01):1250030.

Bancaud, J., Henriksen, O., Rubio-Donnadieu, F., Seino, M., Dreifuss, F. E., and Penry, J. K. (1981). Proposal for revised clinical and electroencephalographic classification of epileptic seizures. *Epilepsia*, 22:489–501.

Bardy, A. H. (1992). Reduction of antiepileptic drug dosage for monitoring epileptic seizures. *Acta Neurol Scand*, 86:466–469.

Bartolomei, F., Wendling, F., Régis, J., Gavaret, M., Guye, M., and Chauvel, P. (2004). Pre-ictal synchronicity in limbic networks of mesial temporal lobe epilepsy. *Epilepsy Research*, 61:89–104.

Bastos, A. M. and Schoffelen, J.-M. (2016). A tutorial review of functional connectivity analysis methods and their interpretational pitfalls. *Frontiers in Systems Neuroscience*, 9.

Baud, M. O., Kleen, J. K., Mirro, E. A., Andrechak, J. C., King-Stephens, D., Chang, E. F., and Rao, V. R. (2018). Multi-day rhythms modulate seizure risk in epilepsy. *Nature Communications*, 9(88):1–10.

Baud, M. O. and Rao, V. R. (2018). Gauging seizure risk. *Neurology*, 91:967–973.

Baud, M. O., Schindler, K., and Rao, V. R. (2021). Under-sampling in epilepsy: limitations of conventional EEG. *Clinical Neurophysiology Practice*, 6:41–49.

Baud, M. O., Vulliemoz, S., and Seeck, M. (2015). Recurrent secondary generalization in frontal lobe epilepsy: Predictors and a potential link to surgical outcome? *Epilepsia*, 56(9):1454–1462.

Bazil, C. W. (2018). Seizure modulation by sleep and sleep state. *Brain Research*, 1703:13–17.

Bazil, C. W. and Walczak, T. S. (1997). Effects of sleep and sleep stage on epileptic and nonepileptic seizures. *Epilepsia*, 38(1):56–62.

Benbir, G., Demiray, D. Y., Delil, S., and Yeni, N. (2013). Interobserver variability of seizure semiology between two neurologist and caregivers. *Seizure*, 22:548–552.

Benignus, V. (1969). Estimation of the coherence spectrum and its confidence interval using the fast fourier transform. *IEEE Transactions on Audio and Electroacoustics*, 17(2):145–150.

Benjamini, Y. and Hochberg, Y. (1995). Controlling the false discovery rate: a practical and powerful approach to multiple testing. *Journal of the Royal Statistical Society: Series B (Methodological)*, 57(1):289–300.

Berg, A. T., Berkovic, S. F., Brodie, M. J., Buchhalter, J., Cross, J. H., Van Emde Boas, W., Engel, J., French, J., Glauser, T. A., Mathern, G. W., Moshé, S. L., Nordli, D., Plouin, P., and Scheffer, I. E. (2010). Revised terminology and concepts for organization of seizures and epilepsies: Report of the ILAE Commission on Classification and Terminology, 2005-2009. *Epilepsia*, 51(4):676–685.

Bernhardt, B. C., Bonilha, L., and Gross, D. W. (2015). Network analysis for a network disorder: The emerging role of graph theory in the study of epilepsy. *Epilepsy & Behavior*, 50:162–170.

Besson, P., Dinkelacker, V., Valabregue, R., Thivard, L., Leclerc, X., Baulac, M., Sammler, D., Colliot, O., Lehéricy, S., Samson, S., and Dupont, S. (2014). Structural connectivity differences in left and right temporal lobe epilepsy. *NeuroImage*, 100:135–144.

Bettus, G., Wendling, F., Guye, M., Valton, L., Régis, J., Chauvel, P., and Bartolomei, F. (2008). Enhanced EEG functional connectivity in mesial temporal lobe epilepsy. *Epilepsy Research*, 81:58–68.

Blum, A. S. and Rutkove, S. B. (2007). *The Clinical Neurophysiology Primer*. Humana Press, Totowa, New Jersey.

Brunet, J. P., Tamayo, P., Golub, T. R., and Mesirov, J. P. (2004). Metagenes and molecular pattern discovery using matrix factorization. *Proceedings of the National Academy of Sciences*, 101(12):4164–4169.

Burns, S. P., Santaniello, S., Yaffe, R. B., Jouny, C. C., Crone, N. E., Bergey, G. K., Anderson, W. S., and Sarma, S. V. (2014). Network dynamics of the brain and influence of the epileptic seizure onset zone. *Proceedings of the National Academy of Sciences*, 111(49):E5321–E5330.

Bénar, C., Chauvière, L., Bartolomei, F., and Wendling, F. (2010). Pitfalls of high-pass filtering for detecting epileptic oscillations: A technical note on "false" ripples. *Clinical Neurophysiology*, 121(3):301–310.

Chapeton, J. I., Haque, R., Wittig, J. H., Inati, S. K., and Zaghloul, K. A. (2019). Large-scale communication in the human brain is rhythmically modulated through alpha coherence. *Current Biology*, 29:2801–2811.

Chauvel, P. and McGonigal, A. (2014). Emergence of semiology in epileptic seizures. *Epilepsy and Behavior*, 38:94–103.

Chen, Z., Grayden, D. B., Burkitt, A. N., Seneviratne, U., D'Souza, W. J., French, C., Karoly, P. J., Dell, K., Leyde, K., Cook, M. J., and Maturana, M. I. (2021). Spatiotemporal patterns of high-frequency activity (80-170 Hz) in long-term intracranial EEG. *Neurology*, 96(7):e1070–e1081.

Cliff, O. M., Lizier, J. T., Tsuchiya, N., and Fulcher, B. D. (2022). Unifying pairwise interactions in complex dynamics. *arXiv*, physics.data-an:2201.11941.

Cole, A. J. (2000). Is epilepsy a progressive disease? The neurobiological consequences of epilepsy. *Epilepsia*, 41(Suppl. 2):S13–S22.

Colominas, M. A., Schlotthauer, G., and Torres, M. E. (2014). Improved complete ensemble EMD: A suitable tool for biomedical signal processing. *Biomedical Signal Processing and Control*, 14:19–29.

Conrad, E. C., Tomlinson, S. B., Wong, J. N., Oechsel, K. F., Shinohara, R. T., Litt, B., Davis, K. A., and Marsh, E. D. (2020). Spatial distribution of interictal spikes fluctuates over time and localizes seizure onset. *Brain*, 143:554–569.

Cook, M. J., Karoly, P. J., Freestone, D. R., Himes, D., Leyde, K., Berkovic, S., O'Brien, T., Grayden, D. B., and Boston, R. (2016). Human focal seizures are characterized by populations of fixed duration and interval. *Epilepsia*, 57(3):359–368.

Cook, M. J., O'Brien, T. J., Berkovic, S. F., Murphy, M., Morokoff, A., Fabinyi, G., D'Souza, W., Yerra, R., Archer, J., Litewka, L., Hosking, S., Lightfoot, P., Ruedebusch, V., Sheffield, W. D., Snyder, D., Leyde, K., and Himes, D. (2013). Prediction of seizure likelihood with a long-term, implanted seizure advisory system in patients with drug-resistant epilepsy: A first-in-man study. *The Lancet Neurology*, 12:563–571.

Cramer, J. A. and French, J. (2001). Quantitative assessment of seizure severity for clinical trials: A review of approaches to seizure components. *Epilepsia*, 42(1):119–129.

Dadok, V. M., Kirsch, H. E., Sleigh, J. W., Lopour, B. A., and Szeri, A. J. (2015). A probabilistic method for determining cortical dynamics during seizures. *Journal of Computational Neuroscience*, 38:559–575.

Davis, K. A., Devries, S. P., Krieger, A., Mihaylova, T., Minecan, D., Litt, B., Wagenaar, J. B., and Stacey, W. C. (2018). The effect of increased intracranial EEG sampling rates in clinical practice. *Clinical Neurophysiology*, 129:360–367.

Davis, K. A., Sturges, B. K., Vite, C. H., Ruedebusch, V., Worrell, G., Gardner, A. B., Leyde, K., Sheffield, W. D., and Litt, B. (2011). A novel implanted device to wirelessly record and analyze continuous intracranial canine EEG. *Epilepsy Research*, 96(1):116–122.

de Cheveigné, A. and Nelken, I. (2019). Filters: When, why, and how (not) to use them. *Neuron*, 102(2):280–293.

de Tisi, J., Bell, G. S., Peacock, J. L., McEvoy, A. W., Harkness, W. F., Sander, J. W., and Duncan, J. S. (2011). The long-term outcome of adult epilepsy surgery, patterns of seizure remission, and relapse: A cohort study. *The Lancet*, 378:1388–1395.

den Heijer, J. M., Otte, W. M., van Diessen, E., van Campen, J. S., Lorraine Hompe, E., Jansen, F. E., Joels, M., Braun, K. P., Sander, J. W., and Zijlmans, M. (2018). The relation between cortisol and functional connectivity in people with and without stress-sensitive epilepsy. *Epilepsia*, 59:179–189.

Dobesberger, J., Ristić, A. J., Walser, G., Kuchukhidze, G., Unterberger, I., Höfler, J., Amann, E., and Trinka, E. (2015). Duration of focal complex, secondarily generalized tonic-clonic, and primarily generalized tonic-clonic seizures - A video-EEG analysis. *Epilepsy and Behavior*, 49:111–117.

Donos, C., Maliia, M. D., Dümpelmann, M., and Schulze-Bonhage, A. (2018). Seizure onset predicts its type. *Epilepsia*, 59(3):650–660.

Dorr, V. L., Caparos, M., Wendling, F., Vignal, J.-P., and Wolf, D. (2007). Extraction of reproducible seizure patterns based on EEG scalp correlations. *Biomedical Signal Processing and Control*, 2:154–162.

Duckrow, R. B. and Spencer, S. S. (1992). Regional coherence and the transfer of ictal activity during seizure onset in the medial temporal lobe. *Electroencephalography and Clinical Neurophysiology*, 82(6):415–422.

Enatsu, R., Jin, K., Elwan, S., Kubota, Y., Piao, Z., O'Connor, T., Horning, K., Burgess, R. C., Bingaman, W., and Nair, D. R. (2012). Correlations between ictal propagation and response to electrical cortical stimulation: A cortico-cortical evoked potential study. *Epilepsy Research*, 101:76–87.

Engel, J. J. and Crandall, P. H. (1983). Falsely localising ictal onsets with depth EEG telemetry during anticonvulsant withdrawal. *Epilepsia*, 24:344–355.

Englot, D. J., Konrad, P. E., and Morgan, V. L. (2016). Regional and global connectivity disturbances in focal epilepsy, related neurocognitive sequelae, and potential mechanistic underpinnings. *Epilepsia*, 57(10):1546–1557.

Ewell, L. A., Liang, L., Armstrong, C., Soltész, I., Leutgeb, S., and Leutgeb, J. K. (2015). Brain state is a major factor in preseizure hippocampal network activity and influences success of seizure intervention. *The Journal of Neuroscience*, 35(47):15635–15648.

Farooque, P. and Duckrow, R. (2014). Subclinical seizures during intracranial EEG recording: Are they clinically significant? *Epilepsy Research*, 108:1790–1796.

Fava, G. A., Park, S. K., and Sonino, N. (2003). Treatment of recurrent depression. *Expert Review of Neurotherapeutics*, 17(15):1109–1117.

Fisher, R. S., Cross, J. H., French, J. A., Higurashi, N., Hirsh, E., Jansen, F. E., Lagae, L., Moshé, S. L., Peltola, J., Perez, E. R., Scheffer, I. E., and Zuberi, S. M. (2017). Operational classification of seizure types by the International League Against Epilepsy: Position paper of the ILAE Commission for Classification and Terminology. *Epilepsia*, 58(4):522–530.

Freestone, D. R., Karoly, P. J., and Cook, M. J. (2017). A forward-looking review of seizure prediction. *Current Opinion in Neurology*, 30:167–173.

Freestone, D. R., Karoly, P. J., Nešić, D., Aram, P., Cook, M. J., and Grayden, D. B. (2014). Estimation of effective connectivity via data-driven neural modeling. *Frontiers in Neuroscience*, 8:1–20.

Freestone, D. R., Kuhlmann, L., Grayden, D. B., Burkitt, A. N., Lai, A., Nelson, T. S., Vogrin, S., Murphy, M., D'Souza, W., Badawy, R., Nesic, D., and Cook, M. J. (2011). Electrical probing of cortical excitability in patients with epilepsy. *Epilepsy and Behavior*, 22:S110–118.

Frusque, G., Borgnat, P., Gonçalves, P., and Jung, J. (2020). Semi-automatic extraction of functional dynamic networks describing patient's epileptic seizures. *Frontiers in Neurology*, 11(579725).

Fulcher, B. D. (2017). Feature-based time-series analysis. *arXiv*, cs.LG:1709.08055.

Galovic, M., de Tisi, J., McEvoy, A. W., Miserocchi, A., Vos, S. B., Borzi, G., Rosillo, J. C., Vuong, K. A., Nachev, P., Duncan, J. S., and Koepp, M. J. (2020). Resective surgery prevents progressive cortical thinning in temporal lobe epilepsy. *Brain*, 143:3262–3272.

Giorgino, T. (2009). Computing and visualizing dynamic time warping alignments in R: The dtw package. *Journal Of Statistical Software*, 31(7):1–24.

Gliske, S. V., Irwin, Z. T., Chestek, C., Hegeman, G. L., Brinkmann, B., Sagher, O., Garton, H. J. L., Worrell, G. A., and Stacey, W. C. (2018). Variability in the location of high frequency oscillations during prolonged intracranial EEG recordings. *Nature Communications*, 9:2155.

Goodfellow, M., Rummel, C., Abela, E., Richardson, M. P., Schindler, K., and Terry, J. R. (2016). Estimation of brain network ictogenicity predicts outcome from epilepsy surgery. *Scientific Reports*, 6:29215.

Gowers, W. R. (1885). *Epilepsy and other chronic convulsive diseases: Their causes, symptoms, and treatment*. William Wood & Company, New York.

Groppe, D. M., Bickel, S., Keller, C. J., Jain, S. K., Hwang, S. T., Harden, C., and Mehta, A. D. (2013). Dominant frequencies of resting human brain activity as measured by the electrocorticogram. *NeuroImage*, 79:223–233.

Guye, M., Régis, J., Tamura, M., Wendling, F., McGonigal, A., Chauvel, P., and Bartolomei, F. (2006). The role of corticothalamic coupling in human temporal lobe epilepsy. *Brain*, 129:1917–1928.

Halford, J. J., Shiau, D., Desrochers, J. A., Kolls, B. J., Dean, B. C., Waters, C. G., Azar, N. J., Haas, K. F., Kutluay, E., Martz, G. U., Sinha, S. R., Kern, R. T., Kelly, K. M., Sackellares, J. C., and LaRoche, S. M. (2015). Inter-rater agreement on identification of electrographic seizures and periodic discharges in ICU EEG recordings. *Clinical Neurophysiology*, 126:1661–1669.

Harden, C. L. and Pennell, P. B. (2013). Neuroendocrine considerations in the treatment of men and women with epilepsy. *The Lancet Neurology*, 12:72–83.

Hodgkin, A. and Huxley, A. (1952). A quantitative description of membrane current and its application to conduction and excitation in nerve. *J Physiol*, 117:500–544.

Howbert, J. J., Patterson, E. E., Stead, S. M., Brinkmann, B., Vasoli, V., Crepeau, D., Vite, C. H., Sturges, B., Ruedebusch, V., Mavoori, J., Leyde, K., Sheffield, W. D., Litt, B., and Worrell, G. A. (2014). Forecasting seizures in dogs with naturally occurring epilepsy. *PLoS ONE*, 9(1):e81920.

Hsu, D., Chen, W., Hsu, M., and Beggs, J. M. (2008). An open hypothesis: Is epilepsy learned, and can it be unlearned? *Epilepsy and Behavior*, 13:511–522.

Huang, N. E., Shen, Z., Long, S. R., Wu, M. C., Shih, H. H., Zheng, Q., Yen, N. C., Tung, C. C., and Liu, H. H. (1998). The empirical mode decomposition and the Hubert spectrum for nonlinear and non-stationary time series analysis. *Proceedings of the Royal Society A: Mathematical, Physical and Engineering Sciences*, 454:903–995.

Izhikevich, E. M. (2003). Simple model of spiking neurons. *IEEE Transactions on Neural Networks*, 14(6):1569–1572.

Janz, D. (1962). The grand mal épilepsies and the sleeping-waking cycle. *Epilepsia*, 3(1):69–109.

Jarosiewicz, B. and Morrell, M. (2021). The RNS system: Brain-responsive neurostimulation for the treatment of epilepsy. *Expert Review of Medical Devices*, 18(2):129–138. PMID: 32936673.

Jiménez-Jiménez, D., Nekkare, R., Flores, L., Chatzidimou, K., Bodi, I., Honavar, M., Mullatti, N., Elwes, R. D. C., Selway, R. P., Valentín, A., and Alarcón, G. (2015). Prognostic value of intracranial seizure onset patterns for surgical outcome of the treatment of epilepsy. *Clinical Neurophysiology*, 126:257–267.

Jirsa, V. K., Stacey, W. C., Quilichini, P. P., Ivanov, A. I., and Bernard, C. (2014). On the nature of seizure dynamics. *Brain*, 137:2210–2230.

Jmail, N., Gavaret, M., Bartolomei, F., and Bénar, C.-G. (2017). Despiking SEEG signals reveals dynamics of gamma band preictal activity. *Physiological Measurement*, 38(2):N42–N56.

Karoly, P. J., Freestone, D. R., Boston, R., Grayden, D. B., Himes, D., Leyde, K., Seneviratne, U., Berkovic, S., O'Brien, T., and Cook, M. J. (2016). Interictal spikes and epileptic seizures: Their relationship and underlying rhythmicity. *Brain*, 139:1066–1078.

Karoly, P. J., Goldenholz, D. M., Freestone, D. R., Moss, R. E., Grayden, D. B., Theodore, W. H., and Cook, M. J. (2018a). Circadian and circaseptan rhythms in human epilepsy: A retrospective cohort study. *The Lancet Neurology*, 17:977–985.

Karoly, P. J., Kuhlmann, L., Soudry, D., Grayden, D. B., Cook, M. J., and Freestone, D. R. (2018b). Seizure pathways: A model-based investigation. *PLoS Computational Biology*, 14(10):e1006403.

Karoly, P. J., Rao, V. R., Gregg, N. M., Worrell, G. A., Bernard, C., Cook, M. J., and Baud, M. O. (2021). Cycles in epilepsy. *Nature Reviews Neurology*, 17(May).

Karoly, P. J., Ung, H., Grayden, D. B., Kuhlmann, L., Leyde, K., Cook, M. J., and Freestone, D. R. (2017). The circadian profile of epilepsy improves seizure forecasting. *Brain*, 140:2169–2182.

Karthick, P., Tanaka, H., Khoo, H., and Gotman, J. (2018). Prediction of secondary generalization from a focal onset seizure in intracerebral EEG. *Clinical Neurophysiology*, 129:1030–1040.

Kaufmann, E., Seethaler, M., Lauseker, M., Fan, M., Vollmar, C., Noachtar, S., and Rémi, J. (2020). Who seizes longest? Impact of clinical and demographic factors. *Epilepsia*, 61:1376–1385.

Khambhati, A. N., Bassett, D. S., Oommen, B. S., Chen, S. H., Lucas, T. H., Davis, K. A., and Litt, B. (2017). Recurring functional interactions predict network architecture of interictal and ictal states in neocortical epilepsy. *eNeuro*, 4(1):e0091–16.2017.

Khambhati, A. N., Davis, K. A., Lucas, T. H., Litt, B., and Bassett, D. S. (2016). Virtual cortical resection reveals push-pull network control preceding seizure evolution. *Neuron*, 91:1170–1182.

Khambhati, A. N., Davis, K. A., Oommen, B. S., Chen, S. H., Lucas, T. H., Litt, B., and Bassett, D. S. (2015). Dynamic network drivers of seizure generation, propagation and termination in human neocortical epilepsy. *PLoS Computational Biology*, 11(12):e1004608.

Kim, D., Cho, J.-W., Lee, J., Joo, E. Y., Hong, S. C., Hong, S. B., and Seo, D.-W. (2011). Seizure duration determined by subdural electrode recordings in adult patients with intractable focal epilepsy. *Journal of Epilepsy Research*, 1(2):57–64.

Kim, H. and Park, H. (2007). Sparse non-negative matrix factorizations via alternating non-negativity-constrained least squares for microarray data analysis. *Bioinformatics*, 23(12):1495–1502.

Kim, J., He, Y., and Park, H. (2014). Algorithms for nonnegative matrix and tensor factorizations: A unified view based on block coordinate descent framework. *Journal of Global Optimization*, 58:285–319.

Kim, J. and Park, H. (2011). Fast nonnegative matrix factorization: An active-set-like method and comparisons. *SIAM Journal on Scientific Computing*, 33(6):3261–3281.

King-Stephens, D., Mirro, E., Weber, P. B., Laxer, K. D., Van Ness, P. C., Salanova, V., Spencer, D. C., Heck, C. N., Goldman, A., Jobst, B., Shields, D. C., Bergey, G. K., Eisenschenk, S., Worrell, G. A., Rossi, M. A., Gross, R. E., Cole, A. J., Sperling, M. R., Nair, D. R., Gwinn, R. P., Park, Y. D., Rutecki, P. A., Fountain, N. B., Wharen, R. E., Hirsch, L. J., Miller, I. O., Barkley, G. L., Edwards, J. C., Geller, E. B., Berg, M. J., Sadler, T. L., Sun, F. T., and Morrell, M. J. (2015). Lateralization of mesial temporal lobe epilepsy with chronic ambulatory electrocorticography. *Epilepsia*, 56(6):959–967.

Kini, L. G., Bernabei, J. M., Mikhail, F., Hadar, P., Shah, P., Khambhati, A. N., Oechsel, K., Archer, R., Boccanfuso, J., Conrad, E., Shinohara, R. T., Stein, J. M., Das, S., Kheder, A., Lucas, T. H., Davis, K. A., Bassett, D. S., and Litt, B. (2019). Virtual resection predicts surgical outcome for drug-resistant epilepsy. *Brain*, 142(12):3892–3905.

Kini, L. G., Davis, K. A., and Wagenaar, J. B. (2016). Data integration: combined imaging and electrophysiology data in the cloud. *Neuroimage*, 124:1175–1181.

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

Kramer, M. A. and Cash, S. S. (2012). Epilepsy as a disorder of cortical network organization. *Neuroscientist*, 18(4):360–372.

Kramer, M. A., Eden, U. T., Cash, S. S., and Kolaczyk, E. D. (2009). Network inference with confidence from multivariate time series. *Physical Review E*, 79(6).

Kramer, M. a., Eden, U. T., Kolaczyk, E. D., Zepeda, R., Eskandar, E. N., and Cash, S. S. (2010). Coalescence and fragmentation of cortical networks during focal seizures. *The Journal of Neuroscience*, 30(30):10076–10085.

Kramer, M. A., Kolaczyk, E. D., and Kirsch, H. E. (2008). Emergent network topology at seizure onset in humans. *Epilepsy Research*, 79:173–186.

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

Kuhnert, M. T., Elger, C. E., and Lehnertz, K. (2010). Long-term variability of global statistical properties of epileptic brain networks. *Chaos*, 20(043126).

Lagarde, S., Buzori, S., Trebuchon, A., Carron, R., Scavarda, D., Milh, M., McGonigal, A., and Bartolomei, F. (2019). The repertoire of seizure onset patterns in human focal epilepsies: Determinants and prognostic values. *Epilepsia*, 60:85–95.

Langdon-Down, M. and Brain, W. R. (1929). Time of day in relation to convulsions in epilepsy. *Lancet*, 213(5516):1029–1032.

Lawn, N. D., Bamlet, W. R., Radhakrishnan, K., O'Brien, P. C., and So, E. L. (2004). Injuries due to seizures in persons with epilepsy. *Neurology*, 63(9):1565–1570.

Le Bouquin-Jeannès, R., Wendling, F., Faucon, G., and Bartolomei, F. (2002). Mise en correspondance de relations inter-structures lors de crises d'épilepsie. *ITBM-RBM*, 23:4–13.

Lee, D. D. and Seung, H. S. (1999). Learning the parts of objects by non-negative matrix factorization. *Nature*, 401:788–791.

Legendre, P. and Gauthier, O. (2014). Statistical methods for temporal and space-time analysis of community composition data. *Proceedings of the Royal Society B: Biological Sciences*, 281:20132728.

Leguia, M. G., Andrzejak, R. G., Rummel, C., Fan, J. M., Mirro, E. A., Tcheng, T. K., Rao, V. R., and Baud, M. O. (2021). Seizure cycles in focal epilepsy. *JAMA Neurology*, 78(4):454–463.

Li, A., Huynh, C., Fitzgerald, Z., Cajigas, I., Brusko, D., Jagid, J., Claudio, A. O., Kanner, A. M., Hopp, J., Chen, S., Haagensen, J., Johnson, E., Anderson, W., Crone, N., Inati, S., Zaghloul, K. A., Bulacio, J., Gonzalez-Martinez, J., and Sarma, S. V. (2021). Neural fragility as an EEG marker of the seizure onset zone. *Nature Neuroscience*, 24(10):1465–1474.

Li, Y. and Ngom, A. (2013). The non-negative matrix factorization toolbox for biological data mining. *Source Code for Biology and Medicine*, 8(1).

Loddenkemper, T., Vendrame, M., Zarowski, M., Gregas, M., Alexopoulos, A. V., Wyllie, E., and Kothare, S. V. (2011). Circadian patterns of pediatric seizures. *Neurology*, 76:145–153.

Louis Door, V., Caparos, M., Wendling, F., Vignal, J.-P., and Wolf, D. (2007). Extraction of reproducible seizure patterns based on EEG scalp correlations. *Biomedical Signal Processing and Control*, 2:154–162.

Luck, S. J. (2014). Basics of Fourier Analysis and Filtering. In *An Introduction to the Event-Related Potential Technique*, chapter 7, pages 219–248. MIT, Cambridge, Massachusetts, 2nd edition.

Macdonald, R. L. and Kelly, K. M. (1995). Antiepileptic drug mechanisms of action. *Epilepsia*, 36(s2):S2–S12.

Mackevicius, E. L., Bahle, A. H., Williams, A. H., Gu, S., Denisenko, N. I., Goldman, M. S., and Fee, M. S. (2019). Unsupervised discovery of temporal sequences in high-dimensional datasets, with applications to neuroscience. *eLife*, 8.

Mantel, N. (1967). The detection of disease clustering and a generalized regression approach. *Cancer Research*, 27(1):209–220.

Marciani, M. G. and Gotman, J. (1986). Effects of drug withdrawal on location of seizure onset. *Epilepsia*, 27(4):423–431.

Mardia, K. V. (1976). Linear-circular correlation coefficients and rhythmometry. *Biometrika*, 63(2):403–405.

Martinet, L. E., Ahmed, O. J., Lepage, K. Q., Cash, S. S., and Kramer, M. A. (2015). Slow spatial recruitment of neocortex during secondarily generalized seizures and its relation to surgical outcome. *Journal of Neuroscience*, 35(25):9477–9490.

Martinet, L.-E., Fiddyment, G., Madsen, J. R., Eskandar, E. N., Truccolo, W., Eden, U. T., Cash, S. S., and Kramer, M. A. (2017). Human seizures couple across spatial scales through travelling wave dynamics. *Nature Communications*, 8:14896.

Maturana, M. I., Meisel, C., Dell, K., Karoly, P. J., D'Souza, W., Grayden, D. B., Burkitt, A. N., Jiruska, P., Kudlacek, J., Hlinka, J., Cook, M. J., Kuhlmann, L., and Freestone, D. R. (2020). Critical slowing down as a biomarker for seizure susceptibility. *Nature Communications*, 11:2172.

McGonigal, A. (2020). Semiology and epileptic networks. *Neurosurgery Clinics of North America*, 31(3):373–385.

Meisel, C., Plenz, D., Schulze-Bonhage, A., and Reichmann, H. (2016). Quantifying antiepileptic drug effects using intrinsic excitability measures. *Epilepsia*, 57(11):e210–e215.

Meisel, C., Schulze-Bonhage, A., Freestone, D., Cook, M. J., Achermann, P., and Plenz, D. (2015). Intrinsic excitability measures track antiepileptic drug action and uncover increasing/decreasing excitability over the wake/sleep cycle. *Proceedings of the National Academy of Sciences*, 112(47):14694–14699.

Mitsis, G. D., Anastasiadou, M. N., Christodoulakis, M., Papathanasiou, E. S., Papacostas, S. S., and Hadjipapas, A. (2020). Functional brain networks of patients with epilepsy exhibit pronounced multiscale periodicities, which correlate with seizure onset. *Human Brain Mapping*, 41:2059–2076.

Myers, C. T. and Mefford, H. C. (2015). Advancing epilepsy genetics in the genomic era. *Genome Medicine*, 7:91.

Naftulin, J. S., Ahmed, O. J., Piantoni, G., Eichenlaub, J. B., Martinet, L. E., Kramer, M. A., and Cash, S. S. (2018). Ictal and preictal power changes outside of the seizure focus correlate with seizure generalization. *Epilepsia*, 59:1398–1409.

Napolitano, C. E. and Orriols, M. A. (2013). Changing patterns of propagation in a super-refractory status of the temporal lobe. Over 900 seizures recorded over nearly one year. *Epilepsy and Behavior Case Reports*, 1:126–131.

Navis, A. and Harden, C. (2016). A treatment approach to catamenial epilepsy. *Current Treatment Options in Neurology*, 18:30.

Nevado-Holgado, A. J., Marten, F., Richardson, M. P., and Terry, J. R. (2012). Characterising the dynamics of EEG waveforms as the path through parameter space of a neural mass model: Application to epilepsy seizure evolution. *NeuroImage*, 59:2374–2392.

Noachtar, S. and Peters, A. S. (2009). Semiology of epileptic seizures: A critical review. *Epilepsy and Behavior*, 15:2–9.

Olsen, R. W. (2018). GABAA receptor: Positive and negative allosteric modulators. *Neuropharmacology*, 136:10–22.

Panagiotopoulou, M., Papasavvas, C. A., Schroeder, G. M., Thomas, R. H., Taylor, P. N., and Wang, Y. (2022). Fluctuations in eeg band power at subject-specific timescales over minutes to days explain changes in seizure evolutions. *Human Brain Mapping*, page 10.1002/hbm.25796.

Park, K. M., Park, S., Lee, D. A., Lee, H.-J., and Kim, S. E. (2020). Longitudinal analysis of structural connectivity in patients with newly diagnosed focal epilepsy of unknown origin. *Clinical Neurology and Neurosurgery*, 199:106264.

Patry, F. L. (1931). The relation of time of day, sleep, and other factors to the incidence of epileptic seizures. *American Journal of Psychiatry*, 87(5):789–813.

Paz, J. T. and Huguenard, J. R. (2015). Microcircuits and their interactions in epilepsy: Is the focus out of focus? *Nature Neuroscience*, 18(3):351–359.

Poldrack, R. A., Laumann, T. O., Koyejo, O., Gregory, B., Hover, A., Chen, M. Y., Gorgolewski, K. J., Luci, J., Joo, S. J., Boyd, R. L., Hunicke-Smith, S., Simpson, Z. B., Caven, T., Sochat, V., Shine, J. M., Gordon, E., Snyder, A. Z., Adeyemo, B., Petersen, S. E., Glahn, D. C., Mckay, D. R., Curran, J. E., Göring, H. H., Carless, M. A., Blangero, J., Dougherty, R., Leemans, A., Handwerker, D. A., Frick, L., Marcotte, E. M., and Mumford, J. A. (2015). Long-term neural and physiological phenotyping of a single human. *Nature Communications*, 6(8885):1–15.

Proix, T., Jirsa, V. K., Bartolomei, F., Guye, M., and Truccolo, W. (2018). Predicting the spatiotemporal diversity of seizure propagation and termination in human focal epilepsy. *Nature Communications*, 9:1088.

Proix, T., Truccolo, W., Leguia, M. G., Tcheng, T. K., King-Stephens, D., Rao, V. R., and Baud, M. O. (2021). Forecasting seizure risk in adults with focal epilepsy: a development and validation study. *The Lancet Neurology*, 20:127–135.

Ramgopal, S., Thome-Souza, S., Jackson, M., Kadish, N. E., Sánchez Fernández, I., Klehm, J., Bosl, W., Reinsberger, C., Schachter, S., and Loddenkemper, T. (2014). Seizure detection, seizure prediction, and closed-loop warning systems in epilepsy. *Epilepsy and Behavior*, 37:291–307.

Ramgopal, S., Thome-Souza, S., and Loddenkemper, T. (2013). Chronopharmacology of anti-convulsive therapy. *Current Neurology and Neuroscience Reports*, 13:339.

Rao, V. R., G. Leguia, M., Tcheng, T. K., and Baud, M. O. (2020). Cues for seizure timing. *Epilepsia*, 62(S1):S15–S31.

Reddy, D. S. and Rogawski, M. A. (2013). Neurosteroids - endogenous regulators of seizure susceptibility and role in the treatment of epilepsy. In Noebels, J. L., Avoli, M., Rogawski, M. A., Olsen, R. W., and Delgado-Escueta, A. V., editors, *Jasper's Basic Mechanisms of the Epilepsies*, pages 984–1002. National Center for Biotechnology Information (US), Bethesda, MD, 4th edition.

Rosenow, F. and Lüders, H. (2001). Presurgical evaluation of epilepsy. *Brain*, 124:1683–1700.

Rossetti, A. O. and Kaplan, P. W. (2010). Seizure semiology: An overview of the 'inverse problem'. *European Neurology*, 63:3–10.

Rummel, C., Goodfellow, M., Gast, H., Hauf, M., Amor, F., Stibal, A., Mariani, L., Wiest, R., and Schindler, K. (2013). A systems-level approach to human epileptic seizures. *Neuroinformatics*, 11:159–173.

Ryzi, M., Brazdil, M., Novak, Z., Hemza, J., Chrastina, J., Oslejskova, H., Rektor, I., and Kuba, R. (2015). Long-term outcomes in patients after epilepsy surgery failure. *Epilepsy Research*, 110:71–77.

Saggio, M. L., Crisp, D., Scott, J., Karoly, P. J., Kuhlmann, L., Nakatani, M., Murai, T., Dümpelmann, M., Schulze-Bonhage, A., Ikeda, A., Cook, M., Gliske, S. V., Lin, J.,

Bernard, C., Jirsa, V., and Stacey, W. (2020). A taxonomy of seizure dynamotypes. *Elife*, 9:e55632.

Sakoe, H. and Seibi, C. (1978). Dynamic programming algorithm optimization for spoken word recognition. *IEEE Transactions on Accoustics, Speech, and Signal Processing*, ASSP-26(1):43–49.

Salami, P., Borzello, M., Kramer, M. A., Westover, M. B., and Cash, S. S. (2021). Quantification of seizure termination patterns reveals limited pathways to seizure end. *medRxiv*, page doi: 10.1101/2021.03.03.21252789.

Salami, P., Peled, N., Nadalin, J. K., Martinet, L. E., Kramer, M. A., Lee, J. W., and Cash, S. S. (2020). Seizure onset location shapes dynamics of initiation. *Clinical Neurophysiology*, 131:1782–1797.

Sammon, J. W. (1969). A nonlinear mapping for data structure analysis. *IEEE Transactions on Computers*, C-18(5):401–409.

Sanchez-Morillo, D., Olaby, O., Fernandez-Granero, M. A., and Leon-Jimenez, A. (2017). Physiological closed-loop control in intelligent oxygen therapy: A review. *Computer Methods and Programs in Biomedicine*, 146:101–108.

Scheiermann, C., Gibbs, J., Ince, L., and Loudon, A. (2018). Clocking in to immunity. *Nature Reviews Immunology*, 18:423–437.

Schevon, C. A., Weiss, S. A., McKhann, G., Goodman, R. R., Yuste, R., Emerson, R. G., and Trevelyan, A. J. (2012). Evidence of an inhibitory restraint of seizure activity in humans. *Nature Communications*, 3:1060.

Schindler, K., Elger, C. E., and Lehnertz, K. (2007a). Increasing synchronization may promote seizure termination: Evidence from status epilepticus. *Clinical Neurophysiology*, 118:1955–1968.

Schindler, K., Gast, H., Stieglitz, L., Stibal, A., Hauf, M., Wiest, R., Mariani, L., and Rummel, C. (2011). Forbidden ordinal patterns of periictal intracranial EEG indicate deterministic dynamics in human epileptic seizures. *Epilepsia*, 52(10):1771–1780.

Schindler, K., Leung, H., Elger, C. E., and Lehnertz, K. (2007b). Assessing seizure dynamics by analysing the correlation structure of multichannel intracranial EEG. *Brain*, 130:65–77.

Schindler, K. A., Bialonski, S., Horstmann, M.-T., Elger, C. E., and Lehnertz, K. (2008). Evolving functional network properties and synchronizability during human epileptic seizures. *Chaos*, 18:033119.

Schroeder, G. M., Chowdhury, F. A., Cook, M. J., Diehl, B., Duncan, J. S., Karoly, P. J., Taylor, P. N., and Wang, Y. (2021). Seizure pathways and seizure durations can vary independently within individual patients with focal epilepsy. *arXiv*, q-bio.NC:2109.06672.

Schroeder, G. M., Diehl, B., Chowdhury, F. A., Duncan, J. S., de Tisi, J., Trevelyan, A. J., Forsyth, R., Jackson, A., Taylor, P. N., and Wang, Y. (2020). Seizure pathways change on circadian and slower timescales in individual patients with focal epilepsy. *Proceedings of the National Academy of Sciences*, 117(20):11048–11058.

Semyanov, A. (2003). Cell type specificity of GABA(A) receptor mediated signaling in the hippocampus. *Curr Drug Targets CNS Neurol Disord*, 2:240–247.

Shah, P., Ashourvan, A., Mikhail, F., Pines, A., Kini, L., Oechsel, K., Das, S. R., Stein, J. M., Shinohara, R. T., Bassett, D. S., Litt, B., and Davis, K. A. (2019). Characterizing the role of the structural connectome in seizure dynamics. *Brain*, 142(7):1955–1972.

Sinha, N., Dauwels, J., Kaiser, M., Cash, S. S., Westover, M. B., Wang, Y., and Taylor, P. N. (2016). Predicting neurosurgical outcomes in focal epilepsy patients using computational modelling. *Brain*, 140(2):319–332.

Sinha, N., Peternell, N., Schroeder, G. M., de Tisi, J., Vos, S., Winston, G., Duncan, J., Wang, Y., and Taylor, P. N. (2021a). Focal to bilateral tonic–clonic seizures are associated with widespread network abnormality in temporal lobe epilepsy. *Epilepsia*, 62.

Sinha, N., Wang, Y., Moreira da Silva, N., Miserocchi, A., McEvoy, A. W., de Tisi, J., Vos, S. B., Winston, G. P., Duncan, J. S., and Taylor, P. N. (2021b). Structural brain network abnormalities and the probability of seizure recurrence after epilepsy surgery. *Neurology*, 96(5):e758–e771.

Sinha, S., Brady, M., Scott, C. A., and Walker, M. C. (2006). Do seizures in patients with refractory epilepsy vary between wakefulness and sleep? *Journal of Neurology, Neurosurgery and Psychiatry*, 77:1076–1078.

Sisterson, N. D., Wozny, T. A., Kokkinos, V., Constantino, A., and Richardson, R. M. (2019). Closed-loop brain stimulation for drug-resistant epilepsy: Towards an evidence-based approach to personalized medicine. *Neurotherapeutics*, 16:119–127.

Spencer, S. S. (2002). Neural networks in human epilepsy: Evidence of and implications for treatment. *Epilepsia*, 43(3):219–227.

Spencer, S. S., Spencer, D. D., Williamson, P. D., and Mattson, R. H. (1981). Ictal effects of anticonvulsant medication withdrawal in epileptic patients. *Epilepsia*, 22:297–307.

Stirling, R. E., Cook, M. J., Grayden, D. B., and Karoly, P. J. (2021). Seizure forecasting and cyclic control of seizures. *Epilepsia*, 62(S1):S2–S14.

Stitt, I., Hollensteiner, K. J., Galindo-Leon, E., Pieper, F., Fiedler, E., Stieglitz, T., Engler, G., Nolte, G., and Engel, A. K. (2017). Dynamic reconfiguration of cortical functional connectivity across brain states. *Scientific Reports*, 7:8797.

Takahashi, H., Takahashi, S., Kanzaki, R., and Kawai, K. (2012). State-dependent precursors of seizures in correlation-based functional networks of electrocorticograms of patients with temporal lobe epilepsy. *Neurological Sciences*, 33:1355–1364.

Taubøll, E., Sveberg, L., and Svalheim, S. (2015). Interactions between hormones and epilepsy. *Seizure*, 28:3–11.

Taylor, P. N., Papasavvas, C. A., Owen, T. W., Schroeder, G. M., Hutchings, F. E., Chowdhury, F. A., Diehl, B., Duncan, J. S., McEvoy, A. W., Miserocchi, A., de Tisi, J., Vos, S. B., Walker, M. C., and Wang, Y. (2021). Normative brain mapping of interictal intracranial EEG to localise epileptogenic tissue. *arXiv*, q-bio.NC:2105.04643.

Thomas, P. J., Olufsen, M., Sepulchre, R., Iglesias, P. A., Ijspeert, A., and Srinivasan, M. (2019). Control theory in biology and medicine: Introduction to the special issue. *Biological Cybernetics*, 113(1):1–6.

Tibshirani, R., Walther, G., and Hastie, T. (2001). Estimating the number of clusters in a data set via the gap statistic. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 63:411–423.

Torres, M. E., Colominas, M. A., Schlotthauer, G., and Flandrin, P. (2011). A complete ensemble empirical mode decomposition with adaptive noise. In *2011 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE.

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

Trimmer, J. S. and Rhodes, K. J. (2004). Localization of voltage-gated ion channels in mammalian brain. *Annual Review of Physiology*, 66:477–519.

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

Ung, H., Baldassano, S. N., Bink, H., Krieger, A. M., Williams, S., Vitale, F., Wu, C., Freestone, D., Nurse, E., Leyde, K., Davis, K. A., Cook, M., and Litt, B. (2017). Intracranial EEG fluctuates over months after implanting electrodes in human brain. *Journal of Neural Engineering*, 14(5).

Ung, H., Davis, K. A., Wulsin, D., Wagenaar, J., Fox, E., McDonnell, J. J., Patterson, N., Vite, C. H., Worrell, G., and Litt, B. (2016). Temporal behavior of seizures and interictal bursts in prolonged intracranial recordings from epileptic canines. *Epilepsia*, 57(12):1949–1957.

van der Maaten, L. V. D., Postma, E., and Herik, J. (2009). Dimensionality reduction: A comparative review. *Technical Report TiCC-TR 2009-005*.

van Mierlo, P., Papadopoulou, M., Carrette, E., Boon, P., Vandenberghe, S., Vonck, K., and Marinazzo, D. (2014). Functional brain connectivity from EEG in epilepsy: Seizure prediction and epileptogenic focus localization. *Progress in Neurobiology*, 121:19–35.

Vaugier, L., McGonigal, A., Lagarde, S., Trébuchon, A., Szurhaj, W., Derambure, P., and Bartolomei, F. (2017). Hyperkinetic motor seizures: A common semiology generated by two different cortical seizure origins. *Epileptic Disorders*, 19(3):362–366.

Vettoretti, M. and Facchinetti, A. (2019). Combining continuous glucose monitoring and insulin pumps to automatically tune the basal insulin infusion in diabetes therapy: A review. *BioMedical Engineering OnLine*, 18(37):1–17.

Vidaurre, D., Smith, S. M., and Woolrich, M. W. (2017). Brain network dynamics are hierarchically organized in time. *Proceedings of the National Academy of Sciences*, 114(48):12827–12832.

Wagenaar, J. B., Brinkmann, B. H., Ives, Z., Worrell, G. A., and Litt, B. (2013). A multimodal platform for cloud-based collaborative research. *International IEEE/EMBS Conference on Neural Engineering*, pages 1386–1389.

Wagner, F. B., Eskandar, E. N., Cosgrove, G. R., Madsen, J. R., Blum, A. S., Potter, N. S., Hochberg, L. R., Cash, S. S., and Truccolo, W. (2015). Microscale spatiotemporal dynamics during neocortical propagation of human focal seizures. *NeuroImage*, 122:114–130.

Wang, Y., Hutchings, F., and Kaiser, M. (2015). Computational modeling of neurostimulation in brain diseases. In *Progress in Brain Research*, volume 222, pages 191–228. Elsevier B.V., 1 edition.

Wang, Y., Schroeder, G. M., Sinha, N., and Taylor, P. N. (2019). Personalised network modelling in epilepsy. *arXiv*, q-bio.NC:1901.01024.

Wang, Y., Trevelyan, A. J., Valentin, A., Alarcon, G., Taylor, P. N., and Kaiser, M. (2017). Mechanisms underlying different onset patterns of focal seizures. *PLoS Computational Biology*, 13(5):e1005475.

Wendling, F., Badier, J., Chauvel, P., and Coatrieux, J. (1997). A method to quantify invariant information in depth-recorded epileptic seizures. *Electroencephalography and Clinical Neurophysiology*, 102:472–485.

Wendling, F., Bartolomei, F., Bellanger, J. J., Bourien, J., and Chauvel, P. (2003). Epileptic fast intracerebral EEG activity: evidence for spatial decorrelation at seizure onset. *Brain*, 126:1449–1459.

Wendling, F., Bartolomei, F., Bellanger, J. J., and Chauvel, P. (2002). Epileptic fast activity can be explained by a model of impaired gabaergic dendritic inhibition. *European Journal of Neuroscience*, 15(9):1499–1508.

Wendling, F., Bellanger, J.-J., Badier, J.-M., and Coatrieux, J.-L. (1996). Extraction of spatio-temporal signatures from depth EEG seizure signals based on objective matching in warped vectorial observations. *IEEE Transactions on Biomedical Engineering*, 43(10):990–1000.

Wendling, F., Shamsollahi, M., Badier, J., and Bellanger, J. (1999). Time-frequency matching of warped depth-EEG seizure observations. *IEEE Transactions on Biomedical Engineering*, 46(5):601–605.

Wenzel, M., Hamm, J. P., Peterka, D. S., and Yuste, R. (2017). Reliable and elastic propagation of cortical seizures in vivo. *Cell Reports*, 19:2681–2693.

Wilson, H. R. and Cowan, J. D. (1972). Excitatory and inhibitory interactions in localized populations of model neurons. *Biophysical journal*, 12:1–24.

Wu, L. and Gotman, J. (1998). Segmentation and classification of EEG during epileptic seizures. *Electroencephalography and Clinical Neurophysiology*, 106:344–356.

Wu, S., Joseph, A., Hammonds, A. S., Celniker, S. E., Yu, B., and Frise, E. (2016). Stability-driven nonnegative matrix factorization to interpret spatial gene expression and build local gene networks. *Proceedings of the National Academy of Sciences*, 113(16):4290–4295.

Zhuang, X., Yang, Z., and Cordes, D. (2020). A technical review of canonical correlation analysis for neuroscience applications. *Human Brain Mapping*, 41:3807–3833.