Slow changes in seizure pathways in individual patients with focal epilepsy

Gabrielle M Schroeder[1], Beate Diehl[2], Fahmida A Chowdhury[2], John S Duncan[2], Jane de Tisi[2], Andrew J Trevelyan[3], Rob Forsyth[3], Andrew Jackson[3], Peter N Taylor[1,2,3], Yujiang Wang[1,2,3]

1. Interdisciplinary Computing and Complex BioSystems Group, School of Computing Science, Newcastle University, UK
2. UCL Queen Square Institute of Neurology, Queen Square, London WC1N 3BG, UK
3. Institute of Neuroscience, Faculty of Medical Science, Newcastle University, UK

Abstract

Personalised medicine requires that treatments adapt to not only the patient, but changing factors within each individual. In focal epilepsy, brain dynamics change over time and modulate pathological processes; however, surprisingly little is known about whether and how seizures vary in the same patient. We quantitatively compared within-subject seizure network dynamics using intracranial recordings of ~700 seizures from 31 patients with focal epilepsy (mean 16.5 seizures/subject) and three canines with focal-onset seizures (mean 62.3 seizures/subject). In all subjects, we found variability in seizure paths through the space of possible network dynamics, producing either a spectrum or clusters of different dynamics. Seizures with similar pathways tended to occur closer together in time, independent of whether antiepileptic medication reduction occurred, but did not necessarily have similar durations or circadian profiles. Our results suggest that slow modulatory processes shape within-subject seizure dynamics, leading to variable seizure pathways that may require tailored treatment approaches.

Focal epilepsy is characterised by spontaneous, recurrent seizures that arise from localised cortical sites[1]. An unresolved question is how much seizure dynamics can vary in individual patients. Past studies suggest that seizures within a single patient share common features[2–6] and progress through a similar sequence[7], or "characteristic pathway[8]," of neural dynamics. However, there is also evidence that seizure dynamics vary in some patients. Clinically, there may be different types of seizure dynamics in patients with multiple seizure onset sites[9], and long-term electroencephalographic (EEG) recordings suggest that a subset of patients have multiple seizure populations with distinct dynamics[8,10–12]. Ictal onset patterns[13,14], the extent of seizure spread[15,16], and seizure recruitment patterns[17] can also differ in the same subject. This variability may arise from fluctuations in the underlying brain state[18–22], suggesting that background neural dynamics affect not only seizure likelihood[19,23], but also seizure *features*. Crucially, a given treatment may only address a subset of a patient's seizure dynamics: for example, a single neurostimulation protocol may not control the complete repertoire of seizures[18] and a single prediction algorithm may fail to forecast all seizures[10,24,25]. 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-subject seizure variability. Is seizure variability present in all patients, and, if so, what form does the variability take? Do within-subject seizures cluster into groups with distinct dynamics? Can other seizure features, such as duration, distinguish different seizure populations[8,10,24]? How are different seizure dynamics distributed in time?

To answer these questions, we need to objectively quantify seizure similarity. 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-subject seizures[26–31], 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. These 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[2,22,32–41], in part due to dynamic changes in the connectivity of the seizure onset zone[7,42–44]. To fully understand how functional interactions support ictal processes, we must also determine if multiple seizure pathways can co-exist in the neural connectivity space of an individual patient. Such

diversity would reveal that the same neural regions can variably interact to produce a variety of pathological dynamics.

Our goal was to quantify and characterise within-subject variability in seizure pathways through network space. We visualised and compared the within-subject seizure network evolutions of human patients with focal epilepsy (recorded for 43-382 hrs) and canine subjects with focal-onset seizures (recorded for 45-475 days). In total, we analysed the network evolutions of 698 seizures (average 16.5 seizures/human subject, 62.3 seizures/canine subject), making our study the first large-scale examination of within-subject seizure variability. In both human and canine recordings, we found variability in seizure network evolution, revealing that within-subject seizures are not well-represented by a single characteristic pathway. However, seizures can also share parts or all of the same pathway, with recurring dynamical elements across seizures. Furthermore, we related variability in seizure network dynamics to seizure duration and the temporal relationships between seizures, providing novel insight into the characteristics of within-subject seizure variability. Our analysis revealed that more similar seizures tend to occur closer together in time in most subjects, suggesting that slow modulatory processes shape seizure pathways.

Results

We analysed seizure network evolution in 31 human subjects (511 seizures total, mean 16.5 seizures/subject) with focal epilepsy who underwent continuous intracranial electroencephalographic (iEEG) recordings as part of presurgical evaluation. Additionally, we analysed seizures from three canine subjects (187 seizures total, mean 62.3 seizures/subject) with naturally occurring epilepsy and focal-onset seizures that underwent chronic (45-475 days) iEEG recordings as part of a seizure prediction study[12,45]. Subject details are provided in Supplementary Tables S1.1 and S1.2.

We first discuss how we visualise seizure network dynamics and quantify the dissimilarity of within-subject seizure pathways through network space. Importantly, differences in seizure network dynamics do not necessarily correspond to anatomical differences in the location and spread of seizure activity; rather, our analysis captures differences in the neural *interactions* that shape ictal processes. We then illustrate key features of this variability using two example subjects and summarise our findings across the entire cohort.

Visualisation and comparison of within-subject seizure network dynamics

We demonstrate our analysis using seven seizures from an example human subject, P1 (for full analysis details, see Methods). From the iEEG recordings of each seizure (Fig. 1a), we computed the sliding-window functional connectivity, defined as band-averaged coherence in six frequency bands (Fig. 1b). Thus, each seizure time window was described by a set of six connectivity matrices that captured interactions between iEEG channels across different frequencies. The set of all possible connectivity patterns creates a high dimensional space, in which each location corresponds to a specific network configuration. As such, each time window was represented by a high-dimensional point, and the evolution of a seizure's network dynamics formed a pathway in this connectivity space. In summary, rather than comparing the seizure iEEG traces, we first transformed the seizures into a network space, thus framing our comparison of seizure dynamics as a comparison of seizure pathways through this feature space.

To visualise seizure pathways through network space, we extracted recurring patterns of seizure connectivity using a dimensionality reduction technique[46,47] (Fig. 1c). Each point in a seizure pathway was thus described as a weighted combination of connectivity "building blocks," each of which corresponded to a specific, recurring seizure network pattern. In our data, at a given time point, a single network pattern often contributed to the majority of the observed seizure connectivity. Therefore, we assigned each time point to the dominant network pattern, resulting in series of network states that provided a simplified description and visualisation of the seizure pathway (see Methods) (Fig. 1d).

In subject P1, we observed four main pathways of network state progressions (Fig. 1d). For the most part, comparing seizures based on these state progressions agreed with our visual impressions of the iEEG traces. However, seemingly similar iEEG traces can be associated with different network structures. This point is illustrated by seizures 1-3: although their dominant ictal activity was in the same spatial location, seizure three was distinguished by a different network state, revealing differing underlying patterns of brain interactions. Meanwhile, iEEG traces with different features can share a common network progression pattern. For example, despite amplitude differences in the ictal discharges, the initial progressions of seizures 4-7 was described by the same state (state 5, green), indicating that a common pattern of brain interactions underlay

early seizure spread in all four seizures. Therefore, these network characterisations of dynamics can reveal hidden seizure features that are not visually accessible in the iEEG traces.

While the state progressions helped visualise differences in seizure pathways, we still lacked an objective quantification of seizure dissimilarity. An ideal measure must compare seizure pathways across three different scenarios, which are all illustrated by subject P1's seizures:

1) Two seizures can progress along the same pathway, but potentially at different rates (e.g., subject P1 seizures 4 and 5).
2) Two seizures can progress along completely distinct pathways (e.g., subject P1 seizures 4 and 3).
3) Two seizures can share portions of the same pathway, but have divergent dynamics during other parts of the seizures (e.g., subject P1 seizures 4 and 7).

We therefore created a measure that recognises similarities in seizure pathways, despite differences in the rates of seizure evolution. After computing a noise-reduced version of the seizure functional connectivity, we applied dynamic time warping[48] 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. Crucially, the warping step ensured that seizures following the same pathway (scenario 1) had a low dissimilarity, regardless of their rates of progression. If two seizures had completely non-overlapping pathways (scenario 2), their dissimilarity would depend on the average distance of the seizures in network space. Finally, if seizures shared part of the same pathway (scenario 3), their dissimilarity was determined by the relative duration of the shared pathway and the distance of the divergent sections of the pathways.

Fig. 1e shows subject P1's seizure dissimilarity matrix, which contains the seizure dissimilarities of all pairs of the subject's seizures. In this subject, the visual agreement between the seizure dissimilarity matrix, the seizure iEEG traces, and the seizure state progressions was striking. Due to their similar pathways, there was a low dissimilarity between seizures 1 and 2, as well as between seizures 4, 5, and 6. Seizure 3 was relatively different from seizures 1 and 2, indicating that their network states occupied distant regions of network space, despite the similarities in their iEEG traces. The seizure dissimilarity matrix also provides a more detailed comparison of seizure dynamics than the simplified state progression representation alone. For example, although

seizures 4-6 all had similar state progressions, seizures 5 and 6 were more similar to each other than to seizure 4 due to subtle network differences that were not captured by the state progression visualisations of the seizure pathways (Supplementary Fig. S2.3).

Therefore, for each subject, the network state progressions provide a simplified description for visualising the seizure network dynamics, while the seizure dissimilarity matrix gives a precise and objective comparison of each pair of seizure pathways. Importantly, both seizure dissimilarity matrices and network state progressions are subject-specific: due to differences in the iEEG implantation, seizure dissimilarities and network states cannot be readily compared across subjects. Throughout the rest of the results, we will focus on the within-subject seizure dynamics of two example subjects, P2 and P3, that highlight important features of within-subject seizure variability, while also summarising findings across the entire cohort. Fig. 2 shows a selection of the example subjects' ictal iEEG traces and the corresponding network state progressions as a reference for the downstream analysis. The seizure variability analysis of all subjects is available on Zenodo (http://dx.doi.org/10.5281/zenodo.3240102) and summarised in Supplementary Table 9.

Seizure dissimilarity matrices quantify differences in within-subject seizure pathways through network space

As in subject P1, the seizure dissimilarity matrices and state progressions revealed variability in seizure pathways in subjects P2 and P3 (Fig. 3a and c). Notably, within each subject there were commonalities in state progressions across seizures, suggesting that seizure dynamics were constrained to certain pathways through network space. However, seizure progression was not always deterministic: in some cases, the same state could lead to seizure termination or further progression along one or more pathways (e.g., subject P2, state 4). A similar flexibility in seizure pathways was observed across subjects.

The seizure dissimilarity matrices quantified these observed differences in the seizure pathways. As expected from the state progressions, there were groups of seizures with similar network progressions and near-zero inter-seizure dissimilarities (e.g., P2 seizures 6-8, P3 seizures 2-4). However, each seizure dissimilarity matrix also revealed lower levels of similarity, such as between seizures 2 and 3 in subject P2, that would be difficult to establish solely from the iEEG traces or state progressions.

We also examined the distribution of seizure dissimilarities in each subject. Strikingly, the seizure dissimilarities in subject P3 (Fig. 3d) had a bimodal distribution, indicating that most pairs of seizures had either relatively similar or different network dynamics, with few intermediate levels of similarity. Meanwhile, subject P2 (Fig. 3b) had a wide range of dissimilarities, suggesting that there were varying degrees of similarity between pairs of seizures in this subject. These different distributions of seizure dissimilarities revealed that seizure variability manifests in different ways across subjects.

Seizures cluster into groups or form a spectrum based on their network dynamics

Given that the distribution of seizure dissimilarities varied across subjects, we asked if within-subject seizures cluster into groups with characteristic network dynamics. Since many subjects, including P2, had intermediate levels of seizure dissimilarity, we first hierarchically clustered each subject's seizures based on their seizure dissimilarity matrix. Rather than assigning seizures to separate groups with different dynamics, the hierarchical clustering described different levels of similarity between seizures. Therefore, to determine if we could group seizures based on their dynamics, we additionally found the optimal number of flat (i.e., non-hierarchical) clusters using the gap statistic, which compares the observed clusters to reference clusters[49] (see Methods). Crucially, the reference distributions also allowed us to test for the absence of multiple seizure clusters. A single seizure cluster in a subject would indicate that 1) all seizures follow the same pathway, forming a single group of seizures with little variability between seizures, or 2) that the seizures form a spectrum of dynamics that is best described by hierarchical relationships, rather than distinct groups of seizures.

The resulting clusters for subjects P2 and P3 are shown in Fig. 4a-b. Although subject P2 had groups of seizures with similar dynamics, the varying levels of similarity between other pairs of seizures meant that there was no optimal way to split these seizures into separate clusters. Instead, the seizures created a spectrum of network dynamics. Meanwhile, the optimal clustering for subject P3 was three seizure clusters, shown in different colours on the dendrogram in Fig. 4b.

Fig. 4c shows the number of seizure clusters across all subjects. The majority of subjects (22 subjects, including 2 canines) had one seizure cluster (i.e., no clear groupings of seizures), 11 subjects (1 canine) had two clusters, and one subject had three clusters. We then examined the distribution of mean seizure dissimilarities in subjects with and without multiple seizure clusters

(Fig. 4d). Although seizure dissimilarities must be compared cautiously across subjects, the mean seizure dissimilarity nonetheless indicates the amount of seizure variability in each subject. We saw a wide range in variability levels in subjects with a single seizure cluster (top histogram): while some had a low average dissimilarity, suggesting that most seizures progress along a similar pathway, others had higher levels of variability, indicating a spectrum of dynamics. However, some seizure variability was present in all subjects, and there was no clear cut-off to distinguish subjects with low and high levels of variability. In subjects with multiple seizure clusters, we observed that there could be variability *within* a seizure cluster (middle histogram), as well as relatively low dissimilarity *between* different seizure clusters (bottom histogram). Thus, the number of seizure clusters does not indicate the level of seizure variability in a given subject, but rather the form of the variability (spectrum vs. clusters).

Differences in seizure temporal duration do not necessarily correspond to differences in seizure network dynamics

Past studies have suggested that seizures with different pathways may be differentiated by their duration; for example, a bimodal distribution of seizure duration would indicate two corresponding groups of seizures with distinct evolutions[8,10,24]. To determine if there is an association between seizure dissimilarities and differences in seizure duration, we created a "duration distance" matrix in each subject that captured the absolute difference in temporal duration between each pair of seizures (Fig. 5b and f) (Methods). In subject P2, the differences between the seizure dissimilarity and duration distance matrices were visually apparent, and there was no association between them (Spearman's $\rho = -0.02$, $p = 0.4866$, one-tailed Mantel test) (Fig. 5a-d). Subject P3, however, had seizure dissimilarity and duration distance matrices with similar structures, and these measures were significantly correlated (Spearman's $\rho = 0.69$, $p = 0.0003$, one-tailed Mantel test) (Fig. 5e-h).

Across subjects (Fig. 5i), Spearman's correlation between seizure dissimilarities and duration distances ranged from -0.29 to 0.86 (mean: 0.33) and was significant in eighteen subjects (52.9%) after global correction for multiple comparisons. In the remaining subjects, there were two possible scenarios that could lower the association between seizure dissimilarity and duration distance: seizures with the same duration could have different network dynamics, or seizures with different durations could have similar network dynamics. Subject P2's seizures demonstrated both of these cases.

We also investigated if the correlation between seizure dissimilarities and duration distances was stronger in subjects that had a clear delineation between short and long seizures (Supplementary S4). However, the existence of clusters in seizure duration was neither necessary nor sufficient for 1) a significant association between seizure dissimilarities and duration distances, or 2) the existence of clusters based on seizure dynamics. As such, seizure duration clusters should be interpreted cautiously, as they may not be associated with differences in seizure pathways.

Seizures with more similar network dynamics tend to occur closer together in time

Many time-varying factors, such as sleep[21,23,50–52] and hormones[53–56], are thought to influence seizure likelihood and dynamics. Additionally, during presurgical monitoring, antiepileptic medication is reduced in many patients, impacting brain dynamics[57]. We therefore explored how seizure variability was distributed in time in each subject. Fig. 6 shows the amount of time elapsed between the seizures of subjects P2 and P3. In both subjects, we saw a shift in the seizure pathways over time. Notably, although subject P3's seizures could be divided into groups based on network dynamics, those seizures were not clustered together in time; instead, there were relatively consistent interictal intervals.

Due to the observed temporal changes in seizure dynamics, we first asked if seizures that occur closer together in time tend to have more similar network dynamics. For each subject, we defined the "temporal distance matrix" as the amount of time elapsed between the onsets of each pair of seizures (Fig. 7b and f). In subject P2, more similar seizures tended to cluster together in time, resulting in a significant correlation between seizure dissimilarities and temporal distances (Spearman's $\rho = 0.69$, $p = 0.001$, one-tailed Mantel test) (Fig. 7 a-d). Meanwhile, subject P3 lacked temporal clusters of similar seizures, and the correlation between seizure dissimilarities and temporal distances was not significant (Spearman's $\rho = 0.24$, $p = 0.0527$, one-tailed Mantel test) (Fig. 7 e-h).

Fig. 7i summarises the relationship between seizure dissimilarities and temporal distances across all subjects. In almost all subjects, there was a positive Spearman's correlation between seizure dissimilarities and temporal distances (range: -0.10 – 0.83, mean: 0.45). This association was significant in 24 subjects (71.0%), including all three canine subjects, demonstrating that the temporal association between similar seizures also exists on longer time-scales. We also explored

whether antiepileptic medication tapering during the presurgical recording was associated with a stronger relationship between seizure dissimilarities and temporal distances (Supplementary S5). Interestingly, there was no association between whether medication tapering was performed and whether the correlation between seizure dissimilarities and temporal distances was significant (χ^2 test, $p = 0.96$), suggesting that other temporal factors also influence seizure dynamics.

Since circadian rhythms influence seizure dynamics in some patients[21,23,50–52], for each subject we also created a "circadian distance matrix" that captured the difference in the time-of-day of the seizures. Only five subjects (14.7%) had significant associations between seizure dissimilarities and circadian distances, indicating that in most subjects, seizure dynamics change on longer time-scales than circadian rhythms (Supplementary S6).

Additionally, we explored if the relationships between seizure dissimilarity and these other seizure features (duration distance, temporal distance, and circadian distance) were highly dependent on the approach used to quantify seizure dissimilarity through network space. Using two alternative measures, we found qualitatively similar results at both the cohort and individual level (Supplementary S10), indicating that the observed associations were robust.

Relationship between seizure variability and clinical factors

Finally, we related seizure variability to clinical factors, including the seizure clinical type, patient surgical outcome, and the pathology of the resected brain tissue. We found that the observed seizure variability was poorly explained by differences in the coarse categorisation of seizure clinical type (subclinical, focal, or secondarily generalised) in most subjects (Supplementary S7). In other words, the observed variability cannot be solely attributed to differences in the symptoms or extent of spread (as defined by the clinical classification) of the seizures. This finding was expected given that seizures of different clinical types can share similar dynamics, while seizures of the same clinical type can have dramatically different features.

We found no association between postsurgical seizure freedom and a number of measures of seizure variability, including the number of seizure clusters, the average seizure dissimilarity, and the number of onset network states (Supplementary S8). These results suggest that the level or form of seizure variability does not impact seizure freedom following surgical resection, perhaps because these measures do not capture the extent or location(s) of the tissue responsible for

generating seizures. Likewise, higher levels of seizure variability were not associated with a particular seizure onset site (Supplementary S8). These findings demonstrate that seizure variability is widely present and suggest that slow temporal factors may be more crucial for determining the extent and form of the variability.

Discussion

We have quantified variability in seizure network dynamics within individual human patients with focal epilepsy, revealing that within-subject seizures are neither deterministic nor comprehensively represented by a single dynamical pathway. Notably, however, in each subject we also observe groups of seizures with shared dynamics, suggesting that seizures are constrained to a subspace of potential brain dynamics. We also find within-subject seizure variability in chronic recordings of three canines, demonstrating that seizure dynamics also vary on longer time-scales. Interestingly, seizure network dynamics change over time in most subjects, with more similar seizures tending to occur closer together in time, suggesting that slow-changing factors modulate within-subject seizure dynamics.

We investigated variability in seizure functional network evolution due to the importance of network interactions in ictal processes[2,7,22,32,34–44] and build on previous work by demonstrating within-subject variability in these pathological network dynamics. However, 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. Meanwhile, comparisons of parameter time courses, derived using model inversion[8,58,59], could reveal different patterns of changes in the neural parameters underlying a patient's seizures. Finally, due to subject-specific recording layouts, we focused on comparing seizure dynamics within individual subjects. However, seizures could also be compared across patients to uncover common classes of pathological dynamics[8,60].

To quantify within-subject variability in seizure network evolution, 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[27,28,30,31] or network[26,29] features computed from scalp or intracranial EEG.

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, the insertion cost increases the dissimilarity of similar seizures with different rates of progression. Although previous work suggested lowering seizure dissimilarity in such scenarios[31], to our knowledge, our dynamic time warping approach provides the first measure of seizure dissimilarity that does not penalise temporal variability between otherwise similar seizures. Despite this difference, these past studies also reported both common and disparate dynamics across within-subject seizures; however, this work 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 almost 700 seizures across thirty-four subjects. Finally, we expand on previous work by using seizure dissimilarity for downstream analysis, including clustering seizures and quantifying the relationship between seizure dynamics and other features.

Previous work has found that within-subject seizures have similar dynamics[2–8], although variability may be introduced through different rates of progression[4,61] or early termination in the seizure pathway[6,8]. In our cohort, we observed that subsets of within-subject 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-subject seizure network dynamics is poorly characterised by a single, characteristic pathway; additionally, seizure variability is not fully described by temporal differences or early termination within the same pathway. We instead propose a model in which various decision points, existing on the framework of potential seizure pathways, produce a repertoire of seizure progressions (Fig. 8). While some parts of the progressions appear deterministic, at other times a decision point may determine 1) the seizure onset state, 2) the next network state, if multiple progressions are possible, or 3) whether the seizure terminates early in the state progression. This model would also explain why seizure variability can either manifest as relatively distinct seizure types or as a spectrum of dynamics. A greater number of decision points, which in turn produce a range of small variations between seizures, would produce a spectrum of seizure dynamics. Fewer decision points and/or separate seizure pathways could produce groups of seizures that each have a characteristic state progression. Importantly, although network states may contain information about pathological tissue[7,43,62,63], the implications of multiple seizure pathways and onset states are uncertain. Further work is needed to determine whether the region responsible for generating seizures and/or its network interactions change across different seizure pathways.

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 progression. We hypothesise that spatiotemporal changes in the interictal neural state produce seizures with different characteristics. Past studies suggest that neural excitability[19,64,65], inhibition[61], and network interactions[22,66] 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[21,50,51], hormones[53–56], and medication[57]. Recently, prolonged recordings of patients with focal epilepsy have revealed that the rates of epileptiform discharges and seizures fluctuate according to both circadian and patient-specific multidien (approximately bi-weekly to monthly) cycles[52]. An intriguing possibility is that the same factors that rhythmically modulate seizure likelihood may also influence seizure dynamics. Surprisingly, although circadian and sleep cycles are known to impact seizure dynamics[21,50,51] and seizure likelihood [23,52], few subjects in our study had seizure variability associated with circadian rhythms, suggesting that factors varying over longer timescales preferentially alter seizure dynamics. Alternatively, arousal levels and sleep stages may be more important than time-of-day in shaping seizure dynamics. Notably, we also observed that seizures with similar state progressions can have different durations in most subjects, suggesting that seizure duration is modulated independently of the seizure pathway. Ultimately, it is likely that various factors, with differential effects on seizure dynamics, interact to produce the observed repertoire of seizure network evolutions.

Notably, a large number of our human subjects 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 dynamics. Changes in antiepileptic medication can impact neural excitability[67–69], and medication tapering increases seizure likelihood in most patients[16,70]; however, it is controversial whether it also affects seizure patterns[9,16,70,71]. In some cases, it appears that medication tapering reveals latent seizure pathways that are suppressed by medication[9] or allows existing pathways to further progress (e.g., the secondary generalisation of typically focal seizures)[16]. 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[64]. Importantly, medication changes alone cannot account for the observed seizure variability in our cohort, as we observed temporal associations of seizure dynamics in patients that did not undergo medication tapering.

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[9] or secondarily generalised[72] 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 dynamics 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[11,12], leading us to underestimate the level of variability in some subjects.

Although seizure variability was not associated with post-surgical seizure freedom, it may have implications for other clinical treatments. For example, in that same patient, seizures with different dynamics may have distinct preictal signatures, making seizure prediction more difficult[10,24]. A successful seizure prediction algorithm would either need to recognise multiple signatures or find common features among the disparate preictal dynamics. Additionally, 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[18]. Thus, such interventions may need to recognise and adapt to the specific characteristics of each seizure type in order to control all seizure dynamics. Importantly, our human 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 shown that there is within-subject variation in seizure network dynamics in subjects with focal epilepsy. This variability is not limited to specific groups of patients, such as those with multifocal seizures; rather, variability in seizure pathways is common across all subjects. We propose that this variability arises from a set of decision points built on a framework of possible seizure progressions. Temporal changes in seizure dynamics suggest that slow-varying factors shape these seizure pathways, perhaps by modulating the background brain state. Further research is needed to determine whether preictal dynamics shape seizure pathways by controlling decisions in the seizure progression. 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.

Methods

Subject and seizure selection: This work was a retrospective study that analysed seizures from 13 patients from the Mayo Clinic and the Hospital of the University of Pennsylvania (available at www.ieeg.org[73,74]) and 18 patients from the University College London Hospital (UCLH) who were diagnosed with refractory focal epilepsy and underwent presurgical monitoring. To explore seizure variability on longer time-scales, intracranial EEG 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[12,45] (available at www.ieeg.org[73,74]). Subjects were selected without reference to cause or other characteristics of their pathology.

For all the iEEG portal patients, all patients gave consent to have their anonymised iEEG data publicly available on the International Epilepsy Electrophysiology Portal (www.ieeg.org)[73,74]. 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).

To be included in the study, each subject 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 subjects, while still ensuring that enough seizures were observed to draw conclusions about the forms, types, and characteristics of seizure variability in each subject. Seizures were excluded from the analysis if they did not have clear electrographic correlates (with clear onset and termination), if they were triggered by or 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 were included in the analysis. Additionally, in the canine subjects, to allow algorithmic identification of seizure termination (see "Seizure extraction in canine subjects"), seizures were only included if there was at least 330 s between the seizure start and the termination of the previous seizure, and if the preictal period (defined as three minutes to one minute before seizure start) lacked large noisy or missing segments.

Additional information about each subject and the analysed seizures is shown in Supplementary Tables S1.1 and S1.2.

Data acquisition: For each human subject, the placement of the intracranial electrodes was determined by the clinical team, independent of this study. In each canine subject, a total of sixteen electrodes, divided into strips of four electrodes, were placed bilaterally on the brain surface[45], again independent of this study. In human subjects, ictal segments were identified and extracted for the analysis based on clinical seizure markings. In canine subjects, seizure start times were previously marked by a team of clinicians, and seizure termination times were determined algorithmically following preprocessing (see "Seizure extraction in canine subjects").

iEEG preprocessing: 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 <0.05 s with the exception of one segment in seizure 2 of subject "Study 020", which was 0.514 s. 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).

Seizure extraction in canine subjects: Because seizure end times were not marked in the canine data, seizure termination was identified algorithmically using an approach similar to Schindler *et al.*[35]. 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. As a reminder, seizures with preictal periods with noisy or missing segments were excluded from the analysis, as were any seizures that occurred within 330 s of the preceding seizure's termination (based on visual inspection).

The absolute slope S of each channel i was given by

$$S_i(t) = |\Delta x_i/\Delta t|$$

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 timepoint 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 5 s moving average. Channel i was considered epileptic at time point t if $S^*_i(t)$ was greater than or equal to 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 1.5 s.

Computing functional connectivity: To compute the time-varying functional connectivity of each seizure, a 10 s sliding window, with 9 s overlap between consecutive windows, was applied to each preprocessed ictal time series. The same sliding window parameters have previously been used to estimate time-varying coherence in ictal iEEG data[75]. 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

$$C_{i,j}(f) = \frac{\left|\sum_{f=f_1}^{f_2} P_{i,j}(f)\right|^2}{\sum_{f=f_1}^{f_2} P_{i,i}(f) \sum_{f=f_1}^{f2} P_{j,j}(f)}$$

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-spectrums and cross-spectrums were calculated using Welch's method (2 s sliding window with 1 s overlap).

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 then 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 patients 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 $6*(n^2 - n)/2$. Each subject's ictal connectivity vectors were subsequently horizontally concatenated together to form a matrix V containing $6*(n^2 - n)/2$ features and m observations, where m is the total number of ictal windows across all seizures.

Non-negative matrix factorization and network state assignment: To extract recurring patterns of functional connectivity and reduce noise in the connectivity matrices, non-negative matrix factorization (NMF)[46] was used to approximately factor each subject's ictal time-varying

connectivity matrix V into two non-negative matrices, W and H, such that $V \approx W \times H$. The matrix W contained subject-specific basis vectors, each of which had $6*(n^2 - n)/2$ features that captured a pattern of connectivity across all channels and frequency bands. Each original ictal time window was summarised 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 factorisations were subject-specific since the basis vector features depended on the iEEG electrode layout in each subject.

To determine the optimal number of basis vectors, r, for each subject, the highest r that produced consistent sets of basis vectors was found (see Supplementary Fig. S2.1 for details). This approach, known as stability NMF[47], exploits the non-deterministic nature of NMF to 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 is consistent with the number of connectivity patterns typically found in ictal iEEG data in other studies[7,42,43]. At each r, NMF of V was performed 25 times using the alternating nonnegative least squares with block principal pivoting method[76,77]. 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.*[47], for each r, the instability I of two sets of basis vectors W and W' was defined as

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

where P is the Pearson's cross-correlation matrix of the sets of basis vectors. 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 $25*(25-1)/2$ pairs of basis vector sets was then averaged to produce $I_{avg}(r)$. The highest r with $I_{avg}(r) \le 0.005$ 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 for the subsequent analysis.

We then used NMF to cluster observations based on the contributions of the basis vectors to each observation[78,79] (Supplementary Fig. S2.2). In our data, most subjects had a sparse coefficients

matrix H, 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 was then described as a progression of network states, enabling visualization of differences in network evolution between seizures.

While the NMF state progressions provided a simplified description of the seizure network dynamics, the entire functional connectivity time courses gave a more accurate description of the dynamics. However, small fluctuations in the connectivity due to noise would create a high baseline dissimilarity between seizures. Therefore, to reduce noise in the connectivity matrices, for each subject the selected factorisation was also used to create $V^*=W\times H$, a lower-rank approximation of the original time-varying seizure functional connectivity (Supplementary S2.2). 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 re-normalised to have an $L1$ norm of 1, ensuring that differences in reconstruction accuracy did not affect the distances between different ictal timepoints.

Computing seizure dissimilarity: Following the NMF-based reconstruction of the seizure connectivity, the network evolution of each seizure was described by a multivariate time series with $6*(n^2 - n)/2$ 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 minimised (Supplementary S3). 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 minimise the $L1$ distance between each pair of seizures, as this metric provides a better measure of distances in high dimensional spaces[80].

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.

We wish to point out as a technical note that due to the warping step, the seizure dissimilarity measure is not a metric distance. Like a metric distance, all dissimilarities are non-negative, the dissimilarity of a seizure to itself is zero, and the dissimilarity between pairs of seizures is symmetric; however, the triangle inequality does not necessarily hold. In particular, any two seizures that follow approximately the same pathway will have a near-zero dissimilarity, regardless of their rates of progression along the pathway. However, their relationship to other seizures that share *part* of the same pathway will depend on how long (temporally) the seizures share the same pathway. Thus, although pairs of seizures may have a low dissimilarity, their relationships to other seizures may differ due to their different rates of progression. These situations can, in turn, lead to violations of the triangle inequality. These limitations should be considered if using the seizure dissimilarity measure as a substitute for a distance measure in future work. In our case, we also compared our dissimilarity measure to two metric distances of trajectories, the Fréchet distance and the Hausdorff distance. Our results are qualitatively similar regardless of the measure used to quantify seizure dissimilarity, and all conclusions still hold (Supplementary S10).

Seizure clustering and cluster evaluation: To identify groups of similar seizures in each subject, each subject's seizures were hierarchically clustered by using the seizure dissimilarity matrix as input for an agglomerative hierarchical clustering algorithm, UPGMA (unweighted pair group method with arithmetic mean). The hierarchical clustering resulted in a dendrogram that summarised the similarity between the subject'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[49], 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 subject. In order to generate reference datasets, the subject's seizures were first projected into Euclidean space using classical (Torgerson's) multidimensional scaling (MDS). 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, …. 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[49].

Comparison to temporal features: To determine if differences in seizure network evolution co-varied with differences in temporal features, three distance matrices were created for each subject:

- temporal distance matrix: the amount of time elapsed (measured in hours) between the onset times of each pair of seizures.
- duration distance matrix: the absolute difference (measured in seconds) in the temporal length of each pair of seizures.
- circadian distance matrix: the difference (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 each subject, Spearman's correlation was computed between the upper triangular elements of the seizure dissimilarity matrix and each of above distance matrices. Since the distances in each matrix were not independent observations, the Mantel test[81] was used to determine the significance of each correlation. Briefly, for each matrix comparison, the rows and columns of one matrix were randomly permuted 10,000 times. The correlation between the two 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 the two matrices. 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 false discovery rate (FDR) correction[82] was applied to the set of p-values from all matrix comparisons across all subjects (34x3 total tests). The correlation was considered significant if the associated adjusted p-value was less than 0.05.

As discussed earlier, the dissimilarity between seizures with partially shared dynamics will partly depend on the temporal duration of the shared dynamics, relative to the warped seizure length. We therefore caution that our seizure dissimilarity measure (computed using dynamic time warping) is not always independent of seizure temporal duration. To determine the robustness of the relationship between seizure dissimilarities and seizure duration distances, as well as the robustness of our other primary results, we additionally computed seizure dissimilarity using two metric distance measures, the Fréchet and Hausdorff distances, which are independent of seizure durations. Using these alternative measures, we then repeated our analysis of seizure clustering and the comparison of seizure dissimilarities with other seizure features (Supplementary S10).

Statistics

The number of seizures analysed in each subject was determined by the number of seizures suitable for analysis (see "Subject and seizure selection") captured during each iEEG recording. These sample sizes are available in Supplementary Tables S1.1 and S1.2. The results focused on qualitative visualisation of within-subject seizure pathways and quantitative comparison of within-subject seizure dynamics, without assigning statistical significance to the similarity of the seizure dynamics. To find an optimal number of seizure clusters based on seizure dynamics in each subject, we used the gap statistic[49] (details in "Seizure clustering and cluster evaluation"). Additionally, in each subject, we used Spearman's correlation to quantify the relationship between the subject's seizure dissimilarity matrix and three distance matrices describing other seizure features (see "Comparison with temporal features"). A *p*-value for each association was then determined using a permutation test (one-tailed Mantel test[81]). Global FDR correction, using the Benjamini-Hochberg algorithm[82], was then applied to all 34x3 (number of subjects x number of within-subject comparisons) *p*-values, and a correlation was considered significant if the associated adjusted p-value was less than 0.05.

Code and data availability

All data was analysed using MATLAB version R2018b. To perform NMF, we used the Nonnegative Matrix Factorization Algorithms Toolbox, available at https://github.com/kimjingu/nonnegfac-matlab/, which implements the alternating nonnegative least squares with block principal pivoting algorithm[76,77]. For the remainder of the analysis, we used MATLAB implementations of standard algorithms (dynamic time warping: `dtw`, hierarchical clustering: `linkage`, multidimensional scaling: `cmdscale`, gap statistic: `evalclusters`, FDR

correction: `mafdr`) or custom code. The iEEG time series of all IEEG Portal subjects is available at www.ieeg.org. The NMF factorisation of each subject's data, along with the code for producing the primary downstream results (state progressions, seizure dissimilarity matrices, clustering, and comparison with temporal features) is published on Zenodo (http://dx.doi.org/10.5281/zenodo.3240102).

Acknowledgements

We thank Gerold Baier, Christoforos Papasavvas, Nishant Sinha, and the rest of the CNNP lab for discussions on the analysis and manuscript. We thank Andrew McEvoy and Anna Miserocchi for undertaking the epilepsy surgery at QS, and Catherine Scott, Roman Rodionov, and Sjoerd Vos for helping with data organisation.

PNT and YW gratefully acknowledge funding from Wellcome Trust (208940/Z/17/Z and 210109/Z/18/Z).

The authors declare no conflict of interest.

Author contributions

G.M.S and Y.W. conceived the idea and developed the methods. B.D. oversaw clinical acquisition and annotation of the UCLH patient EEG data. G.M.S., P.N.T and Y.W. organised the data. G.M.S. performed the visualisation and analysis. Y.W. validated the analysis. G.M.S. drafted the manuscript. All authors participated in critically reviewing and revising the manuscript.

References

1. Rosenow, F. & Lüders, H. Presurgical evaluation of epilepsy. *Brain* **124,** 1683–1700 (2001).
2. Kramer, M. a *et al.* Coalescence and fragmentation of cortical networks during focal seizures. *J. Neurosci.* **30,** 10076–10085 (2010).
3. Schindler, K. *et al.* Forbidden ordinal patterns of periictal intracranial EEG indicate deterministic dynamics in human epileptic seizures. *Epilepsia* **52,** 1771–1780 (2011).
4. Truccolo, W. *et al.* Single-neuron dynamics in human focal epilepsy. *Nat. Neurosci.* **14,** 635–641 (2011).
5. Schevon, C. A. *et al.* Evidence of an inhibitory restraint of seizure activity in humans. *Nat. Commun.* **3,** 1060 (2012).
6. Wagner, F. B. *et al.* Microscale spatiotemporal dynamics during neocortical propagation of human focal seizures. *Neuroimage* **122,** 114–130 (2015).
7. Burns, S. P. *et al.* Network dynamics of the brain and influence of the epileptic seizure onset zone. *Proc. Natl. Acad. Sci.* **111,** E5321–E5330 (2014).
8. Karoly, P. J. *et al.* Seizure pathways : a model-based investigation. *PLoS Comput. Biol.* **14,** e1006403 (2018).
9. Spencer, S. S., Spencer, D. D., Williamson, P. D. & Mattson, R. H. Ictal effects of anticonvulsant medication withdrawal in epileptic patients. *Epilepsia* **22,** 297–307 (1981).
10. Freestone, D. R., Karoly, P. J. & Cook, M. J. A forward-looking review of seizure prediction. *Curr. Opin. Neurol.* **30,** 167–173 (2017).
11. King-Stephens, D. *et al.* Lateralization of mesial temporal lobe epilepsy with chronic ambulatory electrocorticography. *Epilepsia* **56,** 959–967 (2015).
12. Ung, H. *et al.* Temporal behavior of seizures and interictal bursts in prolonged intracranial recordings from epileptic canines. *Epilepsia* **57,** 1949–1957 (2016).
13. Alarcon, G., Binnie, C. D., Elwes, R. D. C. & Polkey, C. E. Power spectrum and intracranial EEG patterns at seizure onset in partial epilepsy. *Electroencephalogr. Clin. Neurophysiol.* **94,** 326–337 (1995).
14. Jiménez-Jiménez, D. *et al.* Prognostic value of intracranial seizure onset patterns for surgical outcome of the treatment of epilepsy. *Clin. Neurophysiol.* **126,** 257–267 (2015).
15. Karthick, P., Tanaka, H., Khoo, H. & Gotman, J. Prediction of secondary generalization from a focal onset seizure in intracerebral EEG. *Clin. Neurophysiol.* **129,** 1030–1040 (2018).
16. Marciani, M. G. & Gotman, J. Effects of drug withdrawal on location of seizure onset. *Epilepsia* **27,** 423–431 (1986).
17. Martinet, L. E., Ahmed, O. J., Lepage, K. Q., Cash, S. S. & Kramer, M. A. Slow spatial recruitment of neocortex during secondarily generalized seizures and its relation to surgical outcome. *J. Neurosci.* **35,** 9477–9490 (2015).
18. Ewell, L. A. *et al.* Brain state is a major factor in preseizure hippocampal network activity and influences success of seizure intervention. *J. Neurosci.* **35,** 15635–15648 (2015).
19. Badawy, R., Macdonell, R., Jackson, G. & Berkovic, S. The peri-ictal state: Cortical excitability changes within 24 h of a seizure. *Brain* **132,** 1013–1021 (2009).
20. Gliske, S. V. *et al.* Variability in the location of high frequency oscillations during prolonged intracranial EEG recordings. *Nat. Commun.* **9,** 2155 (2018).
21. Bazil, C. W. & Walczak, T. S. Effects of sleep and sleep stage on epileptic and nonepileptic seizures. *Epilepsia* **38,** 56–62 (1997).
22. Khambhati, A. N., Davis, K. A., Lucas, T. H., Litt, B. & Bassett, D. S. Virtual cortical resection reveals push-pull network control preceding seizure evolution. *Neuron* **91,** 1170–1182 (2016).
23. Karoly, P. J. *et al.* The circadian profile of epilepsy improves seizure forecasting. *Brain* **140,** 2169–2182 (2017).
24. Cook, M. J. *et al.* Human focal seizures are characterized by populations of fixed duration and interval. *Epilepsia* **57,** 359–368 (2016).
25. Takahashi, H., Takahashi, S., Kanzaki, R. & Kawai, K. State-dependent precursors of seizures in correlation-based functional networks of electrocorticograms of patients with temporal lobe epilepsy. *Neurol. Sci.* **33,** 1355–1364 (2012).

26. Louis Door, V., Caparos, M., Wendling, F., Vignal, J.-P. & Wolf, D. Extraction of reproducible seizure patterns based on EEG scalp correlations. *Biomed. Signal Process. Control* **2,** 154–162 (2007).
27. Wendling, F., Bellanger, J.-J., Badier, J.-M. & Coatrieux, J.-L. Extraction of spatio-temporal signatures from depth EEG seizure signals based on objective matching in warped vectorial observations. *IEEE Trans. Biomed. Eng.* **43,** 990–1000 (1996).
28. Wu, L. & Gotman, J. Segmentation and classification of EEG during epileptic seizures. *Electroencephalogr. Clin. Neurophysiol.* **106,** 344–356 (1998).
29. Le Bouquin-Jeannès, R., Wendling, F., Faucon, G. & Bartolomei, F. Mise en correspondance de relations inter-structures lors de crises d'épilepsie. *ITBM-RBM* **23,** 4–13 (2002).
30. Wendling, F., Shamsollahi, M. B., Badier, J. M. & Bellanger, J. J. Time-frequency matching of warped depth-EEG seizure observations. *IEEE Trans. Biomed. Eng.* **46,** 601–605 (1999).
31. Wendling, F., Badier, J., Chauvel, P. & Coatrieux, J. A method to quantify invariant information in depth-recorded epileptic seizures. *Electroencephalogr. Clin. Neurophysiol.* **102,** 472–485 (1997).
32. Bartolomei, F. *et al.* Pre-ictal synchronicity in limbic networks of mesial temporal lobe epilepsy. *Epilepsy Res.* **61,** 89–104 (2004).
33. Spencer, S. S. Neural networks in human epilepsy : evidence of and implications for treatment. *Epilepsia* **43,** 219–227 (2002).
34. Rummel, C. *et al.* A systems-level approach to human epileptic seizures. *Neuroinformatics* **11,** 159–173 (2013).
35. Schindler, K., Leung, H., Elger, C. E. & Lehnertz, K. Assessing seizure dynamics by analysing the correlation structure of multichannel intracranial EEG. *Brain* **130,** 65–77 (2007).
36. Wendling, F., Bartolomei, F., Bellanger, J. J., Bourien, J. & Chauvel, P. Epileptic fast intracerebral EEG activity: evidence for spatial decorrelation at seizure onset. *Brain* **126,** 1449–1459 (2003).
37. Schindler, K. A., Bialonski, S., Horstmann, M.-T., Elger, C. E. & Lehnertz, K. Evolving functional network properties and synchronizability during human epileptic seizures. *Chaos* **18,** 033119 (2008).
38. Schindler, K., Elger, C. E. & Lehnertz, K. Increasing synchronization may promote seizure termination: evidence from status epilepticus. *Clin. Neurophysiol.* **118,** 1955–1968 (2007).
39. Kramer, M. A. & Cash, S. S. Epilepsy as a disorder of cortical network organization. *Neuroscientist* **18,** 360–372 (2012).
40. Kramer, M. A., Kolaczyk, E. D. & Kirsch, H. E. Emergent network topology at seizure onset in humans. *Epilepsy Res.* **79,** 173–186 (2008).
41. Guye, M. *et al.* The role of corticothalamic coupling in human temporal lobe epilepsy. *Brain* **129,** 1917–1928 (2006).
42. Khambhati, A. N. *et al.* Dynamic network drivers of seizure generation, propagation and termination in human neocortical epilepsy. *PLoS Comput. Biol.* **11,** e1004608 (2015).
43. Khambhati, A. N. *et al.* Recurring functional interactions predict network architecture of interictal and ictal states in neocortical epilepsy. *eNeuro* **4,** e0091–16.2017 (2017).
44. Bettus, G. *et al.* Enhanced EEG functional connectivity in mesial temporal lobe epilepsy. *Epilepsy Res.* **81,** 58–68 (2008).
45. Howbert, J. J. *et al.* Forecasting seizures in dogs with naturally occurring epilepsy. *PLoS One* **9,** e81920 (2014).
46. Lee, D. D. & Seung, H. S. Learning the parts of objects by non-negative matrix factorization. *Nature* **401,** 788–791 (1999).
47. Wu, S. *et al.* Stability-driven nonnegative matrix factorization to interpret spatial gene expression and build local gene networks. *Proc. Natl. Acad. Sci.* **113,** 4290–4295 (2016).
48. Sakoe, H. & Seibi, C. Dynamic programming algorithm optimization for spoken word recognition. *IEEE Trans. Accoustics, Speech, Signal Process.* **ASSP-26,** 43–49 (1978).
49. Tibshirani, R., Walther, G. & Hastie, T. Estimating the number of clusters in a data set via the gap statistic. *J. R. Stat. Soc. Ser. B (Statistical Methodol.* **63,** 411–423 (2001).
50. Bazil, C. W. Seizure modulation by sleep and sleep state. *Brain Res.* **1703,** 13–17 (2018).
51. Sinha, S., Brady, M., Scott, C. A. & Walker, M. C. Do seizures in patients with refractory epilepsy vary between

wakefulness and sleep? *J. Neurol. Neurosurg. Psychiatry* **77,** 1076–1078 (2006).

52. Baud, M. O. *et al.* Multi-day rhythms modulate seizure risk in epilepsy. *Nat. Commun.* **9,** 1–10 (2018).
53. Harden, C. L. & Pennell, P. B. Neuroendocrine considerations in the treatment of men and women with epilepsy. *Lancet Neurol.* **12,** 72–83 (2013).
54. Reddy, D. S. & Rogawski, M. A. Neurosteroids — endogenous regulators of seizure susceptibility and role in the treatment of epilepsy. in *Jasper's Basic Mechanisms of the Epilepsies* (eds. Noebels, J. L., Avoli, M., Rogawski, M. A., Olsen, R. W. & Delgado-Escueta, A. V) 984–1002 (National Center for Biotechnology Information (US), 2013).
55. Taubøll, E., Sveberg, L. & Svalheim, S. Interactions between hormones and epilepsy. *Seizure* **28,** 3–11 (2015).
56. den Heijer, J. M. *et al.* The relation between cortisol and functional connectivity in people with and without stress-sensitive epilepsy. *Epilepsia* **59,** 179–189 (2018).
57. Meisel, C. *et al.* Intrinsic excitability measures track antiepileptic drug action and uncover increasing/decreasing excitability over the wake/sleep cycle. *Proc. Natl. Acad. Sci.* **112,** 14694–14699 (2015).
58. Nevado-Holgado, A. J., Marten, F., Richardson, M. P. & Terry, J. R. 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 (2012).
59. Freestone, D. R. *et al.* Estimation of effective connectivity via data-driven neural modeling. *Front. Neurosci.* **8,** 1–20 (2014).
60. Jirsa, V. K., Stacey, W. C., Quilichini, P. P., Ivanov, A. I. & Bernard, C. On the nature of seizure dynamics. *Brain* **137,** 2210–2230 (2014).
61. Wenzel, M., Hamm, J. P., Peterka, D. S. & Yuste, R. Reliable and elastic propagation of cortical seizures in vivo. *Cell Rep.* **19,** 2681–2693 (2017).
62. Sinha, N. *et al.* Predicting neurosurgical outcomes in focal epilepsy patients using computational modelling. *Brain* **140,** 319–332 (2016).
63. Goodfellow, M. *et al.* Estimation of brain network ictogenicity predicts outcome from epilepsy surgery. *Sci. Rep.* **6,** 29215 (2016).
64. Napolitano, C. E. & Orriols, M. A. Changing patterns of propagation in a super-refractory status of the temporal lobe. Over 900 seizures recorded over nearly one year. *Epilepsy Behav. Case Reports* **1,** 126–131 (2013).
65. Wang, Y. *et al.* Mechanisms underlying different onset patterns of focal seizures. *PLoS Comput. Biol.* **13,** e1005475 (2017).
66. Proix, T., Jirsa, V. K., Bartolomei, F., Guye, M. & Truccolo, W. Predicting the spatiotemporal diversity of seizure propagation and termination in human focal epilepsy. *Nat. Commun.* **9,** 1088 (2018).
67. Meisel, C., Plenz, D., Schulze-Bonhage, A. & Reichmann, H. Quantifying antiepileptic drug effects using intrinsic excitability measures. *Epilepsia* **57,** e210–e215 (2016).
68. Badawy, R. A. B., Macdonell, R. A. L., Berkovic, S. F., Newton, M. R. & Jackson, G. D. Predicting seizure control: cortical excitability and antiepileptic medication. *Ann. Neurol.* **67,** 64–73 (2010).
69. Badawy, R. A. B., Jackson, G. D., Berkovic, S. F. & Macdonell, R. A. L. Cortical excitability and refractory epilepsy: a three-year longitudinal transcranial magnetic stimulation study. *Int. J. Neural Syst.* **23,** 1250030 (2013).
70. Bardy, A. H. Reduction of antiepileptic drug dosage for monitoring epileptic seizures. *Acta Neurol Scand* **86,** 466–469 (1992).
71. Engel, J. J. & Crandall, P. H. Falsely localising ictal onsets with depth EEG telemetry during anticonvulsant withdrawal. *Epilepsia* **24,** 344–355 (1983).
72. Baud, M. O., Vulliemoz, S. & Seeck, M. Recurrent secondary generalization in frontal lobe epilepsy: predictors and a potential link to surgical outcome? *Epilepsia* **56,** 1454–1462 (2015).
73. Wagenaar, J. B., Brinkmann, B. H., Ives, Z., Worrell, G. A. & Litt, B. A multimodal platform for cloud-based collaborative research. *Int. IEEE/EMBS Conf. Neural Eng.* 1386–1389 (2013). doi:10.1109/NER.2013.6696201
74. Kini, L. G., Davis, K. A. & Wagenaar, J. B. Data integration: combined imaging and electrophysiology data in the cloud. *Neuroimage* **124,** 1175–1181 (2016).

75. Martinet, L.-E. *et al.* Human seizures couple across spatial scales through travelling wave dynamics. *Nat. Commun.* **8,** 14896 (2017).

76. Kim, J., He, Y. & Park, H. Algorithms for nonnegative matrix and tensor factorizations: a unified view based on block coordinate descent framework. *J. Glob. Optim.* **58,** 285–319 (2014).

77. Kim, J. & Park, H. Fast nonnegative matrix factorization: an active-set-like method and comparisons. *SIAM J. Sci. Comput.* **33,** 3261–3281 (2011).

78. Brunet, J. P., Tamayo, P., Golub, T. R. & Mesirov, J. P. Metagenes and molecular pattern discovery using matrix factorization. *Proc Natl Acad Sci U S A* **101,** 4164–4169 (2004).

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

80. Aggarwal, C. C., Hinneburg, A. & Keim, D. A. On the surprising behavior of distance metrics in high dimensional space. *Database Theory – ICDT 2001* 420–434 (2001). doi:10.1007/3-540-44503-X_27

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

82. Benjamini, Y. & Hochberg, Y. Controlling the false discovery rate: a practical and powerful approach to multiple testing. *J. R. Stat. Soc. Ser. B* **57,** 289–300 (1995).

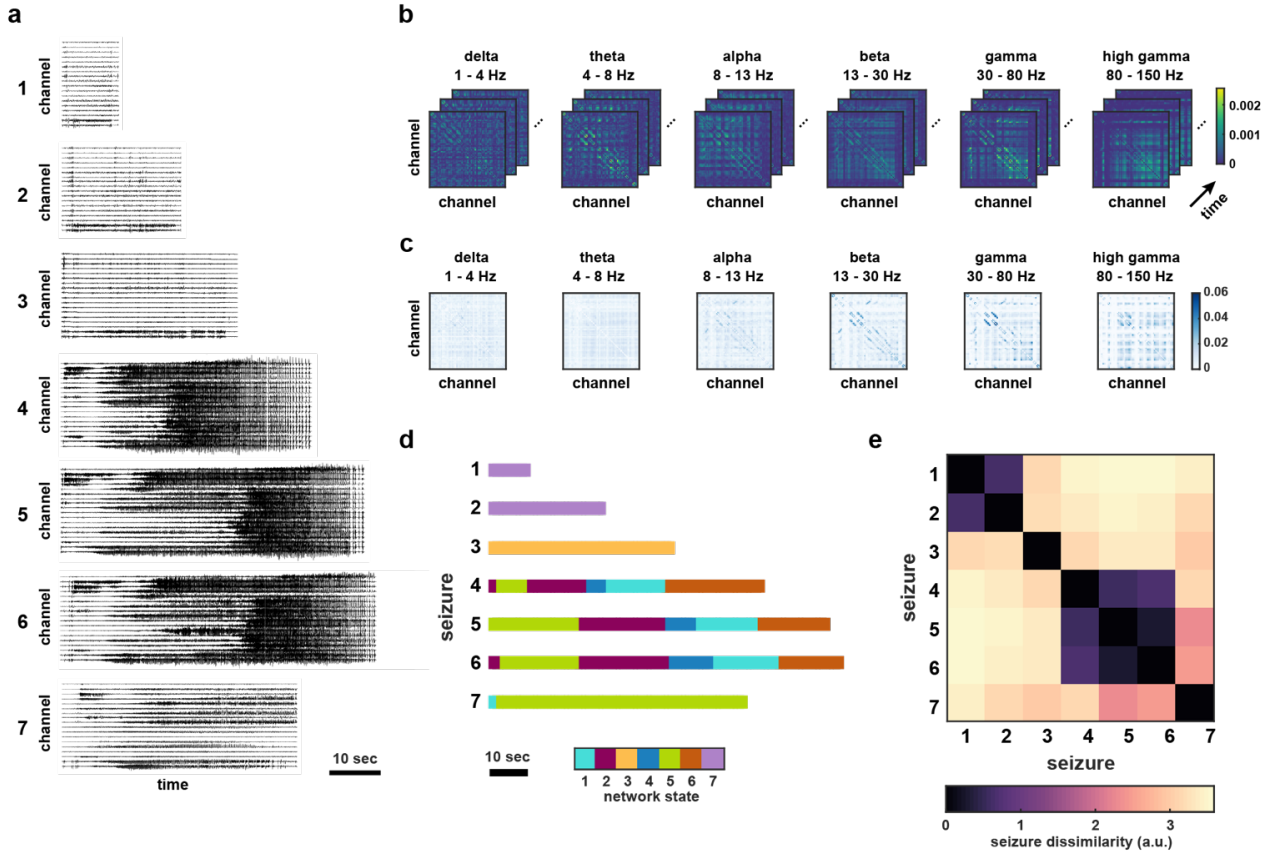

Fig. 1: Visualisation and comparison of seizure network dynamics in an example subject, P1. (a) Intracranial EEG traces of seven seizures of subject P1. (b) The first three windows of the sliding-window functional connectivity, defined as coherence in six different frequency bands, of seizure 4. The entire network evolution of the seizure was described by six sets of connectivity matrices. Each connectivity matrix was normalised such that the upper triangular elements sum to one. (c) Example seizure network state (state 5), derived using non-negative matrix factorisation. (d) State progressions of subject P1's seizures, which provide a visual summary of the seizure's pathway through network space. Each state is indicated by a different colour. (e) Subject P1's seizure dissimilarity matrix, 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.

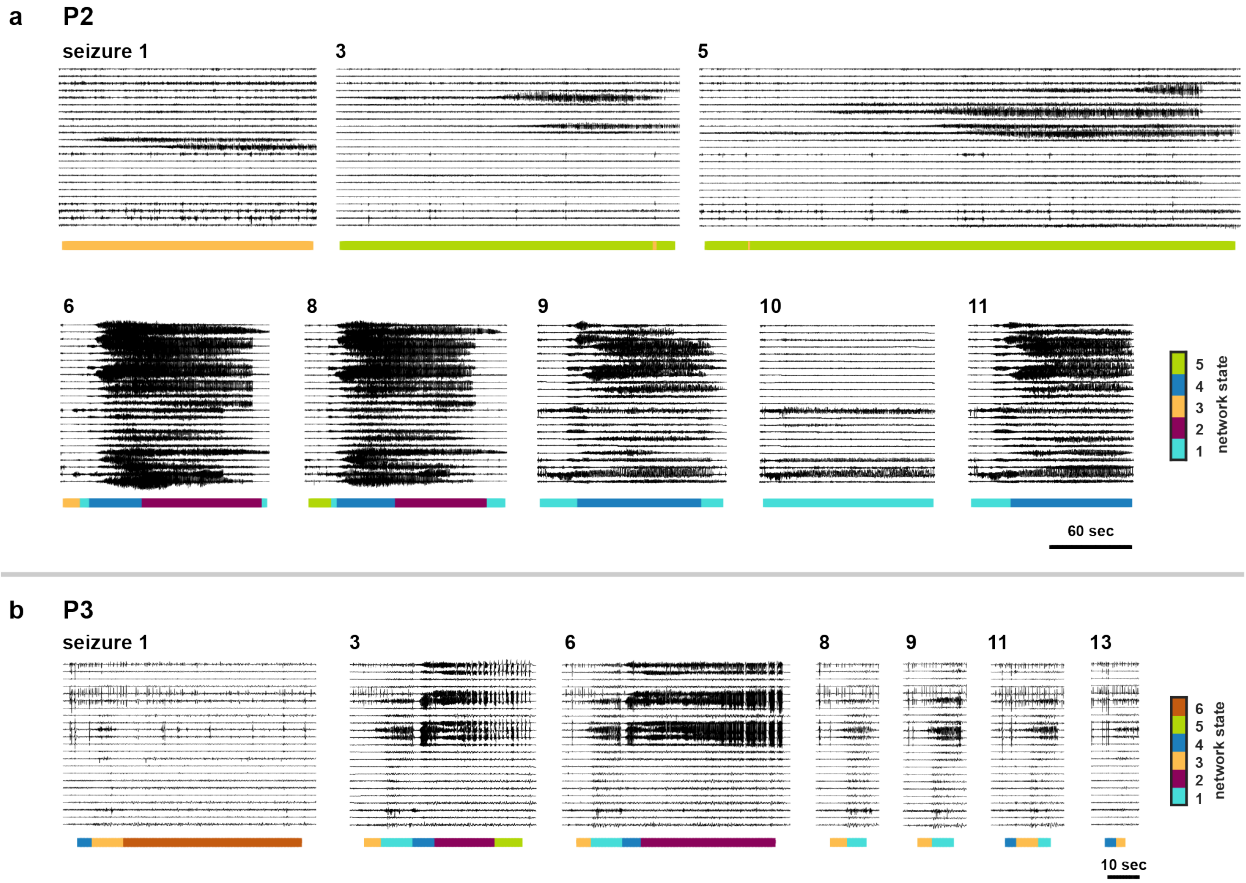

Fig. 2: Within-subject variability in seizure network progressions. The iEEG traces and corresponding network states of selected seizures from subjects P2 (a) and P3 (b). Seizures are numbered by the order of their occurrence in each subject. The seizure network progressions of subjects P2 and P3 were described by five (a) and six (b) network states, respectively, with each network state indicated by a different colour. The state progression of each seizure is placed beneath the ictal iEEG, with each state in the centre of the corresponding time window. Note that due to the 10 s sliding window, each network state corresponds to 10 s of the iEEG trace; thus, transitions in the dynamics seen on the iEEG may not be exactly aligned to changes in the network states.

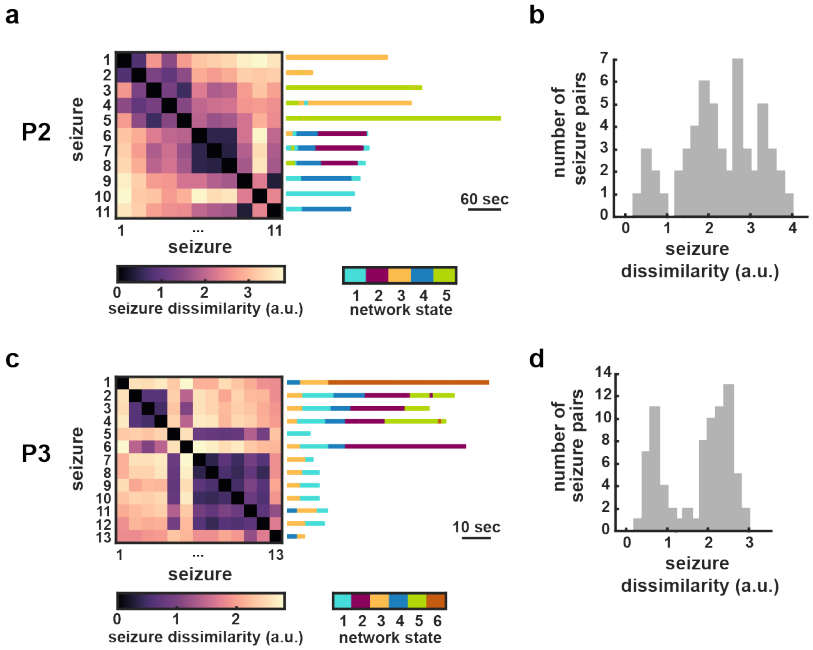

Fig. 3: Seizure dissimilarity matrices and distributions in example subjects. The seizure dissimilarity matrices of subjects P2 (a) and P3 (c) describe the dissimilarity in the network evolution of each pair of seizures. A low dissimilarity (close to zero) indicates that the two seizures have similar network evolutions. To the right of each matrix, the corresponding state progressions of each seizure are shown, allowing a comparison between seizure dissimilarities and network state progressions. For example, in subject P2 (a), there were low dissimilarities between seizures 6-8, all of which had similar network progressions. (b, d) The distributions of seizure dissimilarities in each subject. Note that in both histograms, each observation corresponds to a seizure pair, rather than a single seizure. Subject P2 (b) had a wide range of seizure dissimilarities, while in subject P3 (d), there were either relatively low or high dissimilarities between seizures, forming a bimodal distribution.

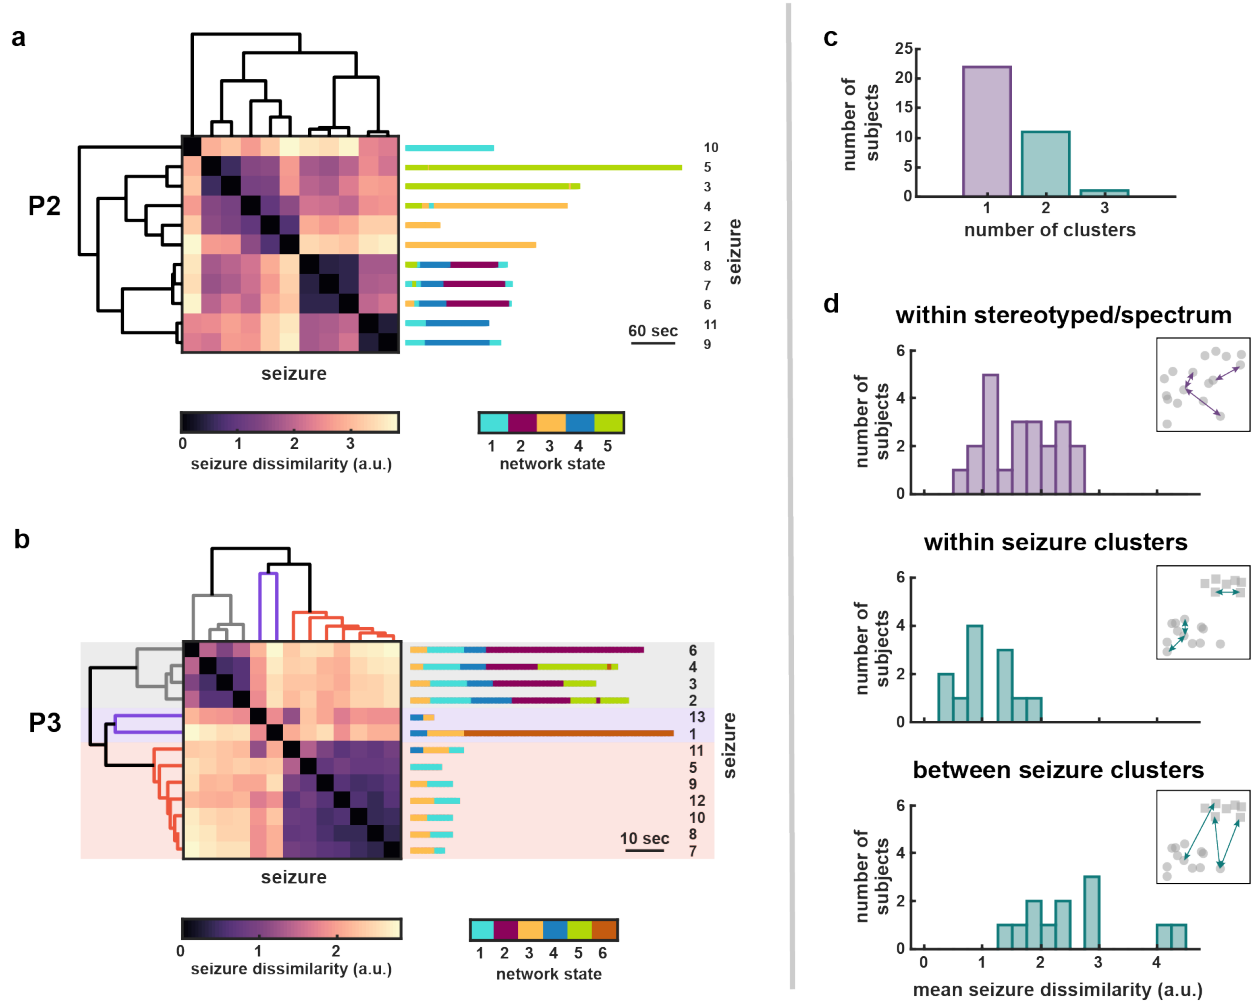

Fig. 4: The form and amount of seizure variability differs across subjects. Seizure clustering results of subjects P2 (a) and P3 (b). The seizure dissimilarity matrices and seizure state progressions are the same as in Fig. 3, but now sorted to match the seizure order of the dendrograms, which describe the hierarchical clustering of the seizures. 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. (a) In subject P2, an optimal non-hierarchical clustering of seizures was not found; instead, seizures were best described by the hierarchical clustering. (b) In subject P3, seizure dynamics were best described by three non-hierarchical clusters (grey, purple, and red dendrogram leaves). (c, d) Analysis of seizure variability across subjects. Histograms and bars in purple correspond to data from subjects with a single seizure cluster, while those in teal correspond to data from subjects with two or more seizure clusters. (c) Bar chart of the number of seizure clusters in each subject. (d) Histograms of seizure dissimilarities, averaged across pairs of seizures within a subject. The top histogram shows dissimilarities in subjects with a single seizure cluster,

963 indicating that dynamics were either stereotyped or formed a spectrum. The bottom two
964 histograms show seizure dissimilarities within and between seizure clusters in the remaining
965 subjects, each of which had at least two seizure clusters. The inset of each histogram shows a
966 schematic illustration of the type of variability (spectrum vs. clustered) and the type of distance
967 (within vs. between cluster/spectrum) investigated. Each gray point represents a seizure, and
968 arrows between seizures provide examples of the dissimilarities used in the computation.

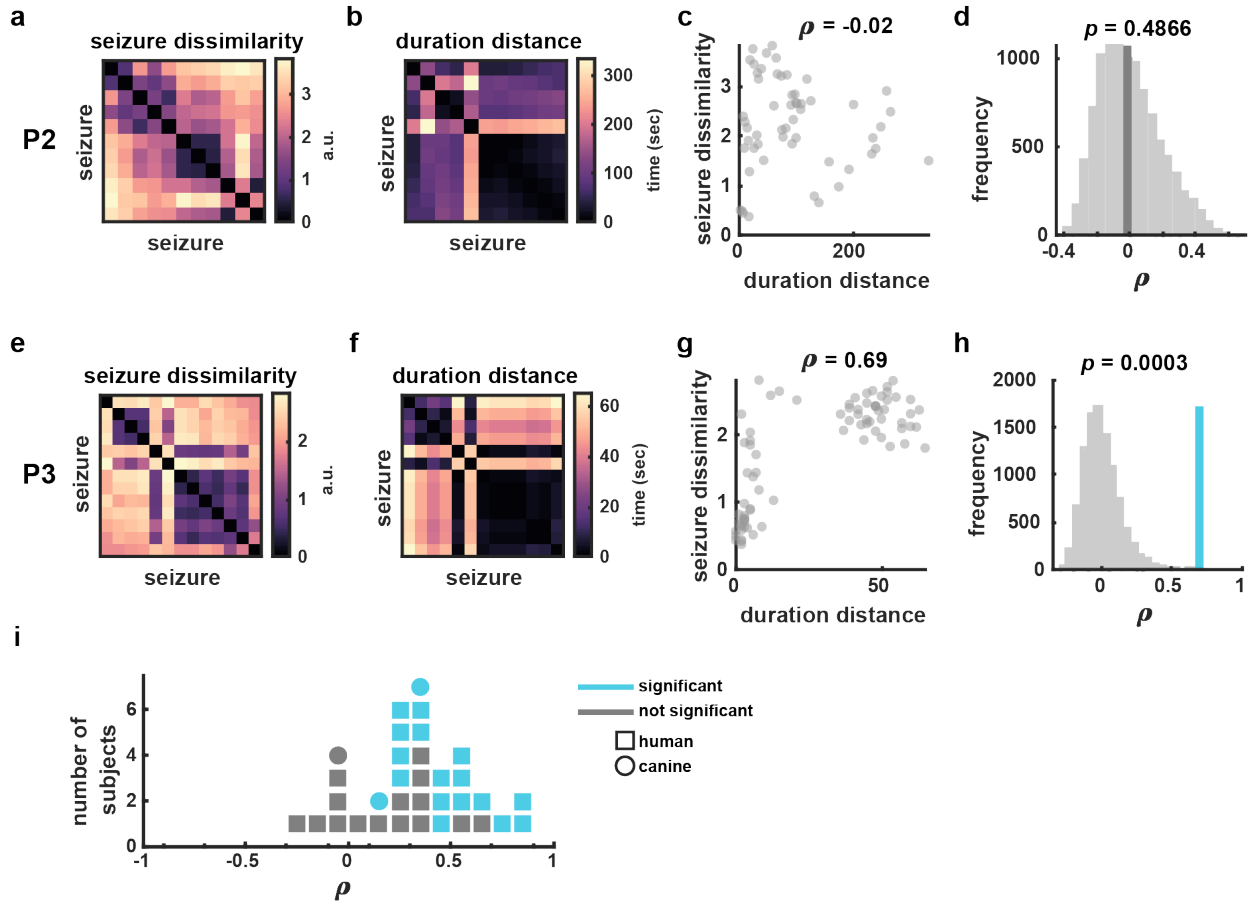

Fig. 5: Comparison of seizure dissimilarities and duration distances. The comparison of seizure dissimilarities and duration distances is shown for subjects P2 (a-d) and P3 (e-h), along with results across all subjects (i). (a and e) Seizure dissimilarity matrices, summarising differences in seizure network dynamics within each subject (same as Fig. 3). (b, f) Duration distance matrices. Each entry corresponds to the absolute difference in seizure duration, in seconds, between two seizures. (c, g) Scatter plots of seizure dissimilarities vs. duration distances, along with the Spearman correlation, ρ, between the two measures. (d, h) For each subject, permutation tests yielded a distribution of 10,000 correlation values that described the expected correlation if there were no relationship between seizure dissimilarities and duration distances. The *p*-value of the association was equal to the proportion of times a correlation value greater than or equal to the observed correlation (vertical bar) was seen in the distribution. The colour of the vertical bar indicates whether the association between seizure dissimilarities and duration distances was significant (blue = significant, grey = not significant after false discovery rate correction). (i) Dot plot showing the range of correlations between seizure dissimilarities and duration distances across

984 all subjects. Each marker represents a subject (square = human subject, circle = canine subject,
985 blue = significant, grey = not significant after false discovery rate correction).

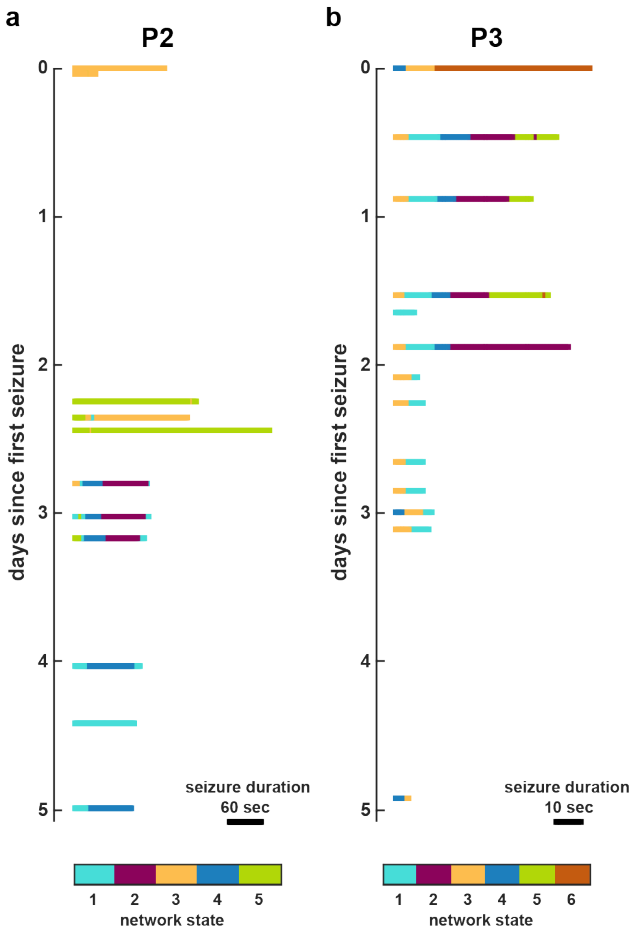

Fig. 6: Temporal distribution of seizure dynamics in example subjects. The seizure state progressions of subjects P2 (a) and P3 (b) are plotted in the order of their occurrence. The vertical distance between seizure progressions is proportional to the amount of time elapsed between seizures.

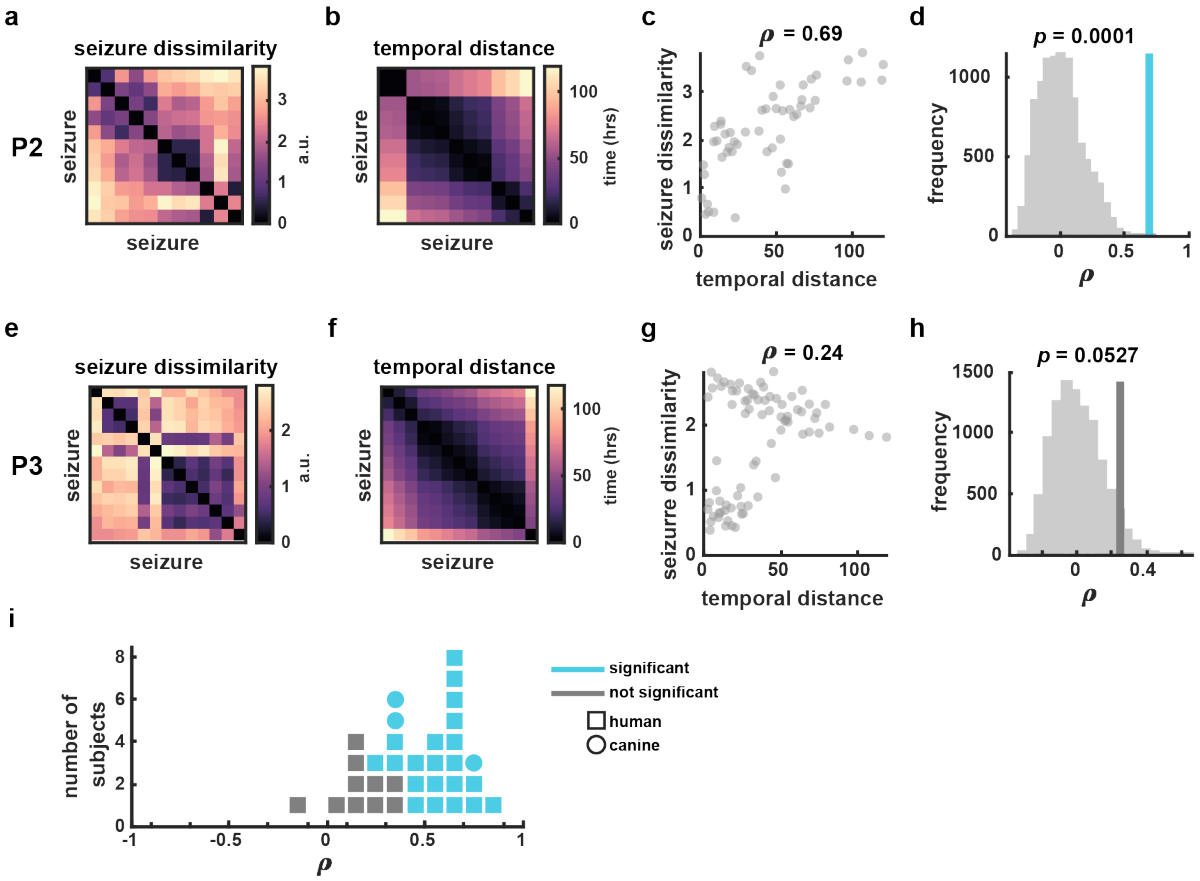

Fig. 7: Comparison of seizure dissimilarities and temporal distances. The comparison of seizure dissimilarities and temporal distances is shown for subjects P2 (a-d) and P3 (e-h), along with results across all subjects (i). Colour and marker coding is the same as in Fig. 5. (a, e) Seizure dissimilarity matrices, summarising differences in seizure network dynamics within each subject (same as in Fig. 3 and 5). (b, f) Temporal distance matrices. Each entry corresponds to the amount of time elapsed between the onsets of a pair of seizures. (c, g) Scatter plots of seizure dissimilarities vs. temporal distances, along with the Spearman correlation, ρ, between the two measures. (d and h) Permutation test results for each subject. See Fig. 5 for a description of the permutation test and *p*-values. (i) Dot plot showing the range of correlations between seizure dissimilarities and duration distances across all subjects, as well as whether each relationship was significant after false discovery rate correction.

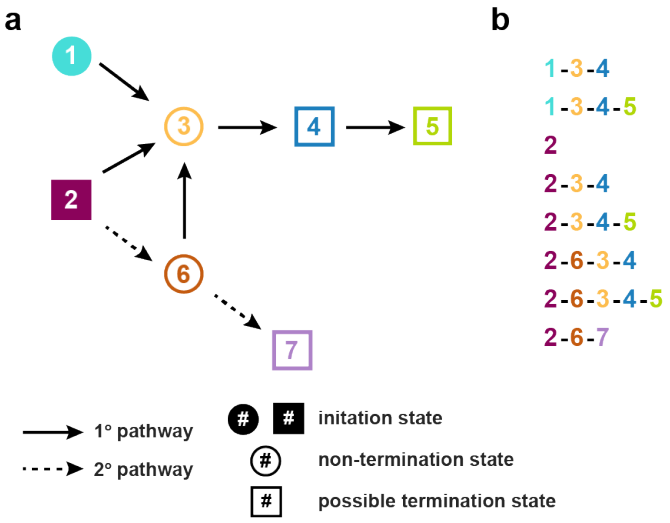

Fig. 8. Hypothesised model for variability in seizure pathways. (a) Diagram of possible seizure pathways, which are described as transitions between seven network states. For simplicity, we use a schematic of seizure progression that provides examples of seizure variability features observed in our data. States that are filled in (states 1 and 2) are possible initiation states in the seizure pathway. Dotted arrows represent secondary transitions that are less likely to occur. 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, which can both lead to state 3), 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) Potential seizures arising from these seizure pathways, demonstrating variability in state onset, state progression, state termination, and state inclusion. All these types of variability are observed in our cohort. Note that the last three progressions, beginning with the state sequence (2, 6), will be rarer since these transitions are less likely.