

A pairwise maximum entropy model uncovers the white matter scaffold underlying emergent dynamics in intracranial EEG

Arian Ashourvan^{1,3,†,*}, Preya Shah^{1,3}, Adam Pines¹, Shi Gu⁵, Christopher W. Lynn⁶, Danielle S. Bassett^{1,2,4,6,7}, Kathryn A. Davis^{1,3}, and Brian Litt^{1,3,4}

¹Department of Bioengineering, School of Engineering and Applied Science, University of Pennsylvania

²Department of Electrical and Systems Engineering, School of Engineering and Applied Science, University of Pennsylvania

³Penn Center for Neuroengineering and Therapeutics, University of Pennsylvania

⁴Department of Neurology, Hospital of the University of Pennsylvania

⁵Department of Computer Science and Technology, University of Electronic Science and Technology of China

⁶Department of Physics & Astronomy, School of Arts & Sciences, University of Pennsylvania

⁷Department of Psychiatry, Perelman School of Medicine, University of Pennsylvania

[†]Correspondence to: Arian Ashourvan, 210 S. 33rd Street, Philadelphia, PA 19104-6321,

Email: ashourv@seas.upenn.edu

ABSTRACT

A major challenge in systems neuroscience is to understand how the brain's structural architecture gives rise to its complex functional dynamics. Here, we address this challenge by examining the inter-ictal activity of five patients with medically refractory epilepsy during ~15 hours of multi-channel intracranial recording. By constructing a pairwise maximum entropy model (MEM) of the observed neural dynamics, we seek to uncover the fundamental relationship between functional activity and its underlying structural substrate. Despite only incorporating the pairwise correlations in the observed neural activity, we find that the pairwise MEM robustly fits large-scale patterns of inter-ictal power dynamics across a wide range of frequency bands, notably displaying time-invariance and cross-frequency similarity. Furthermore, across all frequency bands, we demonstrate that the pairwise MEM accurately identifies the structural white matter connections between brain regions, outperforming other common model-free measures of functional connectivity. Together, our findings show that a simple pairwise MEM, which is explicitly ignorant of higher-order correlations between three or more brain regions, not only captures complex spatiotemporal patterns of neural activity across the frequency spectrum, but also suggests that the network of structural connections in the human brain is a plausible scaffold capable of supporting observed wide-band neural dynamics.

Keywords: Pairwise maximum entropy model; Intracranial EEG (iEEG); Structural connectivity; Structure-function coupling

Introduction

An age-old question in neuroscience concerns the enigmatic relationship between the brain's large-scale function and its underlying white matter structure (Bassett et al., 2018; Bassett and Sporns, 2017). How does the complex array of functional dynamics that we observe in the brain emerge from the static architecture of structural connections between neurons and brain regions? This question lies at the heart of several ongoing efforts, from understanding the brain's cognitive capabilities and development, to localizing and treating human disease. Previous studies demonstrate that network models of functional connectivity (FC), calculated over several minutes of resting state fMRI, share important organizational features with network models of the underlying white-matter structural connectivity (SC) (Skudlarski et al., 2008; Hagmann et al., 2008; Honey et al., 2009; Hermundstad et al., 2013, 2014; Van den Heuvel and Sporns, 2013; Wang et al., 2015). Additionally, recent research shows that accounting for indirect and higher-order structural connections between brain regions can further improve predictions of the FC that is supported (Goñi et al., 2014; Becker et al., 2018). Coupling between FC and SC is not constant over time, but rather evolves throughout adolescence (Hagmann et al., 2010), and is notably altered in disorders of mental health (Skudlarski et al., 2010; Zhang et al., 2011). Moreover, studies have discovered hierarchical and small-world organization in both the functional connectivity between brain regions as well as in the network of structural wires connecting them (Meunier et al., 2009b, 2010; Bassett et al., 2008; Bassett and Bullmore, 2006), and the dynamics of these functional and structural networks have been implicated in cognition (Braun et al., 2015; Bassett et al., 2011; Shine et al., 2016; Bassett et al., 2015), development (Supekar et al., 2009; Dosenbach et al., 2010; Meunier et al., 2009a), and psychiatric conditions (Micheloyannis et al., 2006; Menon, 2011; Yerys et al., 2017).

Commonly used measures of FC, such as the Pearson correlation, can reveal emergent patterns of neural activity at short time scales (e.g., Ashourvan et al. (2017)) and can also provide an accurate mapping of the brain's functional organization at longer time scales (Kramer et al., 2011, 2010). Such model-free measures of FC, however, fundamentally quantify the functional similarity between neural units (e.g., neurons or brain regions), and therefore are limited in their ability to discern between different underlying mechanisms (Adachi et al., 2011; Lu et al., 2011). For example, given two brain regions with a high Pearson correlation, one cannot distinguish between the following three scenarios: (i) The regions are communicating directly via a structural connection, (ii) the regions are communicating indirectly via a structural pathway bridging intermediate regions, or (iii) the two regions are not communicating at all, but are rather being driven by a common third region. To distinguish between direct and indirect communication, and between pairwise versus higher-order interactions, one must begin with a model of how patterns of activity are generated in the brain, and then infer the network of underlying interactions that best describes correlations in the data.

One model of neural activity – the pairwise maximum entropy model (MEM) – has generated particular interest, primarily because it is formally minimal in the sense that it accounts for the observed pairwise correlations in the data while remaining explicitly ignorant with regard to all higher-order correlations (Jaynes, 1957). In recent studies, this model has proven sensitive to the spatiotemporal patterns of co-activation in neuronal spiking at the micro-scale (Schneidman et al., 2006; Shlens et al., 2006; Yeh et al., 2010) and in patterns of blood-oxygen-level-dependent (BOLD) activity at the macro-scale (Roudi et al., 2009; Watanabe et al., 2013, 2014; Ashourvan et al., 2017). The pairwise MEM is based on the principle of maximum entropy, which states that the probability distribution (i.e., the model) that best represents one's current state of knowledge about a system is the one with the largest *entropy* (or *uncertainty*) in the context of previously observed data. Fitting a pairwise MEM entails iteratively adjusting the individual activation of regions as well as the interaction strength of all region pairs until the estimated covariance matrix matches the covariance observed in the data. In this way, the MEM makes quantitative predictions about the frequencies of global activity patterns, rather than simply quantifying the similarities between regions, as is common in studies of FC. A good fit of the pairwise MEM indicates that the observed patterns of power amplitude states can be accurately described with a combination of first-order activations and second-order (pairwise) co-activations between electrodes. A poor fit signifies that higher-order effects (e.g. interactions between triplets or quadruplets of regions) and/or common inputs from other external regions may be required to explain the behavior of the system. Notably, prior studies using resting state fMRI data have demonstrated that the pairwise MEM accurately predicts the observed patterns of regional activations, and provides a more accurate map of the underlying structural connectivity (derived from diffusion tensor imaging (DTI)) than conventional FC-based methods (Watanabe et al., 2013).

Here, we build on prior work by examining intracranial EEG (iEEG) power dynamics across several frequency bands, ranging between 4–180 Hz, using a pairwise MEM as our measure of FC. We consider iEEG data acquired from five patients with medically refractory epilepsy over 14.6 ± 2.5 hours (in one-hour segments) of inter-ictal recording. Unlike scalp EEG, which suffers from noise and selective high frequency filtering (Grech et al., 2008), iEEG provides a unique opportunity to record relatively focal neural activity with high temporal resolution from distributed brain structures across a wide range of frequencies. We define 'on'-'off' activation states for each brain region by binarizing the normalized envelope of the power amplitude of each frequency band (Figure 1). Our use of the power amplitude envelope to define the activation states is motivated by several factors. First, it has been shown that the BOLD fMRI signal echoes the envelope of high frequency neural

activity (Logothetis et al., 2001; Winder et al., 2017), specifically when measured by iEEG (Lachaux et al., 2007; Jacques et al., 2016; Nir et al., 2008; Ojemann et al., 2013; Betzel et al., 2018; Reddy et al., 2018). Thus, in light of studies linking BOLD fMRI FC to white matter SC (Honey et al., 2009; Watanabe et al., 2013), we hypothesize that the power of iEEG recordings should also exhibit a clear relationship with the underlying SC. Second, a large body of evidence has demonstrated that iEEG oscillations play an important role in healthy cognitive function, and that breakdowns in oscillatory power are linked to cognitive disorders (Fries, 2009; Wang, 2010; Uhlhaas and Singer, 2006). Third and finally, by defining our activity states using band-passed power amplitudes, we gain the ability to analyze patterns of activity across both time and frequency, thereby avoiding inconsistent phase relationships between electrodes caused by spatially non-stationary signal sources, such as iEEG spiral waves (Huang et al., 2010).

Our results reveal that across all patients, the estimated interaction matrices from the pairwise MEM are more consistent across both time and frequency than the corresponding FC matrices defined simply by pairwise correlations. These results suggest that the intrinsic interactions between brain regions may remain relatively stationary over time, and that a common underlying mechanism may support functional dynamics across all frequency bands. We hypothesize that the physical wiring of SC likely comprises this common scaffold, and we test this hypothesis by comparing the network topology of the inferred maximum entropy interactions with the architecture of structural connections between brain regions. Receiver operator curves (ROCs) reveal that the pairwise MEM robustly identifies the presence or absence of structural connections between regions across all frequency bands for all patients. Furthermore, although the FC defined by power amplitude correlations in the θ and α bands allow identification of the underlying SC in a few patients, the pairwise MEM generates a consistently stronger prediction of SC across all frequency bands and all patients. Despite the similarity between the empirically observed covariance structure and that predicted by the pairwise MEM, we find that empirical observations of large-scale iEEG states often diverge from MEM predictions, notably in γ bands. These observations suggest that higher-order interactions and/or external sources may contribute to the observed activation states. Despite this divergence in the γ bands, our findings generally suggest that regional activation and pairwise co-activation, supported by an underlying scaffold of structural wiring, can explain a large portion of the power amplitude patterns across a wide range of frequencies (4–180 Hz). These results offer new avenues for clinical interventions based on targeted alteration of a patient's structural connectome, with the ultimate goal of influencing neurophysiological dynamics while preserving the cortical real estate and minimizing damage to SC as documented in more aggressive surgical options (such as callosotomy (Concha et al., 2006; Song et al., 2003; Concha et al., 2007; Yogarajah et al., 2010; Taylor et al., 2018)).

Results

The pairwise MEM exhibits high cross-frequency similarity and temporal stationarity

Functional connectivity representing correlations in iEEG states has been shown to display patterns that vary significantly across both frequency and time (Chu et al., 2012; Kramer et al., 2010, 2011). By contrast, given our hypothesis that the inferred maximum entropy parameters should reflect the underlying SC, one would expect these pairwise interactions to maintain a similar architecture across frequency bands. Indeed, the matrix J_{ij} , which encodes the inferred pairwise interactions between electrodes, displays greater cross-frequency similarity than the corresponding FC matrices representing Pearson correlations in power amplitude. Specifically, we found that 4 out of the 5 subjects displayed significantly higher average pairwise correlation between J_{ij} matrices across frequency bands than that of power amplitude FC matrices (t -test, $p = 6.30 \times 10^{-3}$, $p = 6.53 \times 10^{-4}$, $p = 0.78$, $p = 1.39 \times 10^{-6}$, $p = 2.54 \times 10^{-9}$, respectively; SI-Figure 1). Furthermore, the average interaction strengths in the J_{ij} matrices did not differ significantly across most frequency bands (t -test, FDR $p < 0.05$); the notable exception is the high γ band, wherein the average interaction strength was significantly larger than in other frequencies bands. Interestingly, the relation between the estimated pairwise interactions and the Euclidean distances between electrodes was consistent across frequency bands, while an analogous relation between power amplitude FC and Euclidean distance was only observed in the low frequency band (see SI-Figure 2). Consistent with prior studies demonstrating that iEEG FC is relatively stable when estimated over large time windows (Chu et al., 2012), we found that both the pairwise maximum entropy interactions and the power amplitude FC, when estimated over 1 hour intervals, were consistent across time. However, a statistical comparison between the pairwise correlations of FC matrices (over all possible one-hour segments) and that of J_{ij} matrices, revealed that the pairwise MEM interactions were significantly more stable across all patients and frequency bands (t -test, $p < 0.05$; SI-Figure 3). Together, these results support the notion that brain regions' intrinsic propensity for interactions are similar across a wide range of frequencies and display time-invariance over several hours.

The pairwise MEM estimated from iEEG predicts structural connectivity estimated from diffusion imaging

In contrast to common FC measured based on Pearson correlations, the pairwise MEM attempts to explain patterns of activity in the data by identifying underlying interactions between pairs of regions that can account for the observed correlations. Because the pairwise MEM takes into account the indirect propagation of influence between the electrodes, we hypothesized

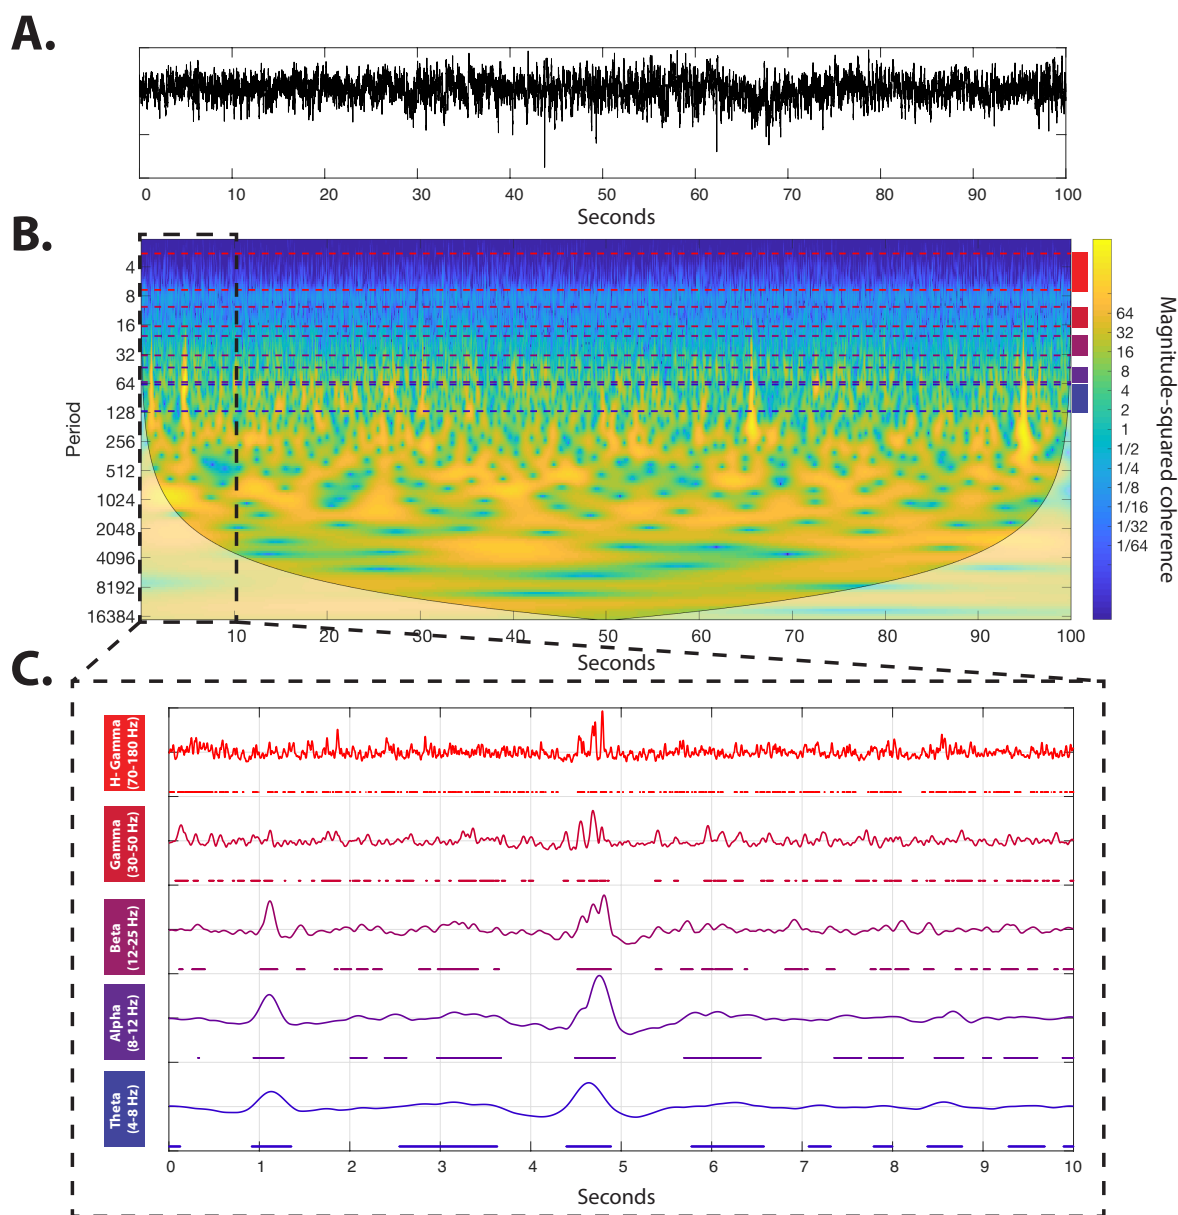


Figure 1. Construction of power amplitude states from iEEG time-series. (A) The raw iEEG time-series recorded from a sample electrode. (B) A wavelet decomposition of 100 seconds of raw iEEG time-series (Grinsted et al., 2004). Color-coded dashed lines show the θ , α , β , γ , and high γ frequency bands. (C) The average power of all frequencies within each band is high-pass filtered at 0.5 Hz, and normalized to obtain a z-score. The resultant time series were then thresholded (>0) separately for each band, to create the binarized power amplitude states on which the MEM operates. Note that the 'on' states are marked by color-coded dots under each curve.

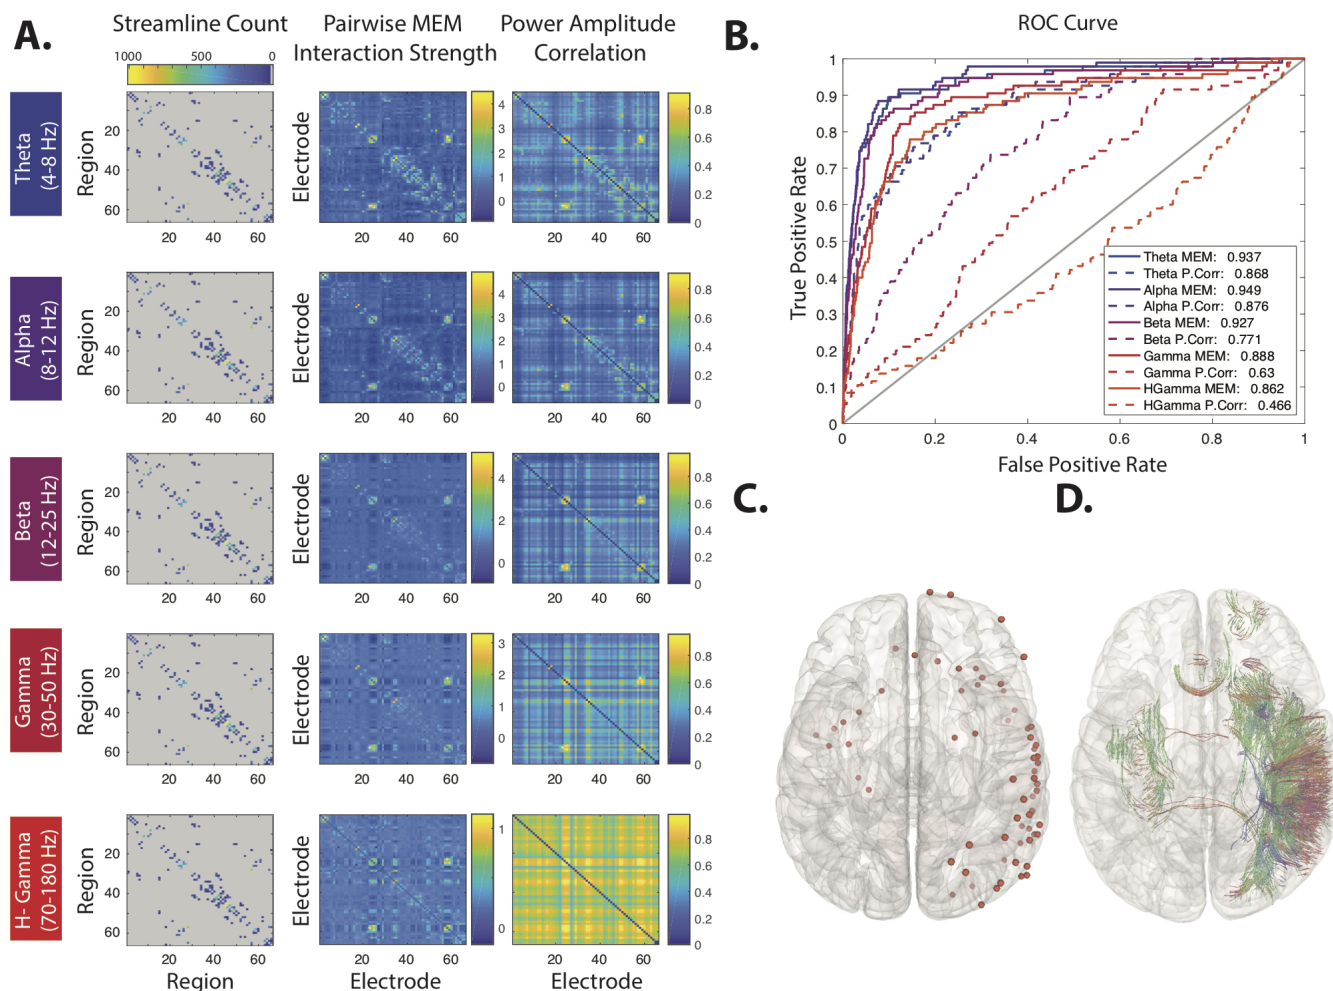


Figure 2. Predicting structural connectivity from estimates of functional connectivity. (A) (Left Column) A structural connectivity matrix from a single representative patient (#5), each element of which indicates the number of streamlines estimated between a pair of brain regions. (Middle Column) The interaction matrices, J_{ij} , obtained from the pairwise MEM. (Right Column) The functional connectivity matrices estimated from band-passed power amplitude correlations. For regions covered by multiple electrodes, we selected only the electrode closest to the region's centroid for display in these plots as well as in the ROC analysis. (B) ROC curves for identification of anatomically connected regions based on interaction matrices (solid lines) and band-passed power amplitude correlations (dashed lines). Lines are color-coded based on the frequency band. (C) Brain overlay providing the position of the selected electrodes (red dots). (D) Streamlines connecting the regions covered by the selected electrodes in the same patient. Here the 3-dimensional orientations of the streamlines are coded by color.

that the estimated maximum entropy interaction matrices would be sparser than the FC matrices obtained from power amplitude correlations. Moreover, we hypothesized that the pairwise MEM would identify underlying white-matter tracts that support the observed functional activity. Consistent with our first hypothesis, we find that the J_{ij} matrices are notably sparser than the FC matrices obtained from the power amplitude correlations, and display a distribution of strengths that is more skewed (see SI-Figure 2). Additionally, consistent with our second hypothesis, we find that the J_{ij} matrices are highly predictive of the underlying (binarized) SC across all frequency bands, as determined by a receiver-operator curve (ROC) analysis (see Figure 2.A–D for results obtained in a single representative subject, and see Figure 3 for consistent results across all subjects). The ability of the pairwise MEM to identify the underlying SC is even more pronounced when noting that the FC matrices, obtained from both power amplitude correlations and traditional time series correlations, are significantly less predictive of white matter SC (t -test, FDR $p < 0.05$; SI-Figure 4) across all frequencies, patients, and SC binarization thresholds (SI-Figure 5). Furthermore, across all frequency bands, 4 out of the 5 subjects displayed a significant positive relationship between the weights of the SC matrices and the weights of the interaction matrices obtained from the pairwise MEM (SI-Figure 6). Together, these results provide converging evidence that the pairwise MEM robustly uncovers the underlying structural connectivity between regions and, conversely, that this structural wiring provides a scaffold that supports the observed wide-band cortical dynamics.

Accounting for distance in the relation between the pairwise MEM and structural connectivity

Extensive prior work has demonstrated that the weights of structural connections between brain regions decreases as a function of inter-regional distance (Donahue et al., 2016; Kaiser and Hilgetag, 2004; Lewis et al., 2008; Rubinov et al., 2015). The effects of distance on SC are even more pronounced with respect to iEEG data because the electrode grids and strips are distributed focally rather than uniformly across the brain. Indeed, we observe that the most structurally connected electrode pairs are proximal and that long range fibers are scarce (SI-Figure 7.A). In fact, the distance between electrodes is highly predictive of the presence of a structural connection between them (average AUC across subject, $= 0.986 \pm 0.003$). In order to better account for the role of distance, we repeated the ROC analysis after removing a variable percentage of the closest electrode pairs (SI-Figure 7.B). With the exception of one subject's data in the θ and α bands, we observed that the interaction matrices estimated by the pairwise MEM continued to provide more accurate predictions of SC than the FC matrices obtained from power amplitude correlations. In a complementary approach, we regressed out the inter-regional Euclidean distances from both the maximum entropy interaction strengths as well as the power amplitude correlations. We then repeated the ROC analysis, finding, across all frequency bands, that the FC matrices obtained from power amplitude correlations did not significantly predict the presence of SC (SI-Figure 7.C). By contrast, this distance effect is markedly weaker for the interaction matrices obtained from the pairwise MEM, with the exception of the high- γ band in 4 out of the 5 subjects. Together, these results suggest that, although the presence of SC is highly confounded with inter-regional distance, especially in the high- γ band, the pairwise MEM is able to detect direct structural connections between brain regions beyond those anticipated by inter-regional distance alone.

The pairwise MEM robustly predicts the observed covariance structure of power amplitude states

Because a large portion of the cortex is canvassed by electrodes in our patient sample, we have the necessary coverage to measure interactions between distant brain regions. However, we expect to see some degree of divergence between the observed and predicted state probabilities for three reasons: (i) increasing the dimensionality of the system in general decreases the accuracy of a fit due to an exponential increase in the state space, (ii) unmeasured brain regions cannot be included in the model and are therefore likely to induce systematic errors, and (iii) real-world neural activity is likely to exhibit some degree of higher-order interactions. We assessed the performance of the pairwise MEM first by comparing the empirically observed covariance matrices with the computationally estimated covariance matrices obtained from the model. We estimated the model-based covariance matrices by randomly sampling the probability distribution using the Metropolis-Hastings algorithm (see Methods). We found that the estimated covariance matrices are statistically similar to the empirical covariance matrices across all frequency bands (average correlation values between empirical and estimated covariance matrices for all patients, 0.96 ± 0.01 , 0.95 ± 0.07 , 0.94 ± 0.03 , 0.98 ± 0.01 , 0.96 ± 0.04 , $p < 0.001$ Bonferroni correction, respectively). Differences between the estimated and empirical covariance matrices existed along the diagonal, where estimated values are significantly larger than empirically measured values (all patients and frequency bands, paired t -test, $p < 0.001$ Bonferroni correction, see SI-Figure 8), and also among pairs of electrodes located near one another physically (as seen in Figure 4). We also observed a rapid decrease in estimation errors as a function of the distance between electrodes across all frequency bands (see Figure 4.E for results obtained from a single subject, and see SI-Figure 9 for results obtained from all subjects). The similar and highly significant ($p \approx 0$) exponential fit to the covariance estimation error versus electrode pair distance across all subjects and frequency bands highlights the much higher estimation error for nearby electrode pairs (SI-Figure 9). In summary, these results demonstrate that the pairwise MEM provides a statistically robust prediction of the observed covariance structure of the power amplitude states across all frequency bands, and that its relatively small approximation errors are predominantly local.

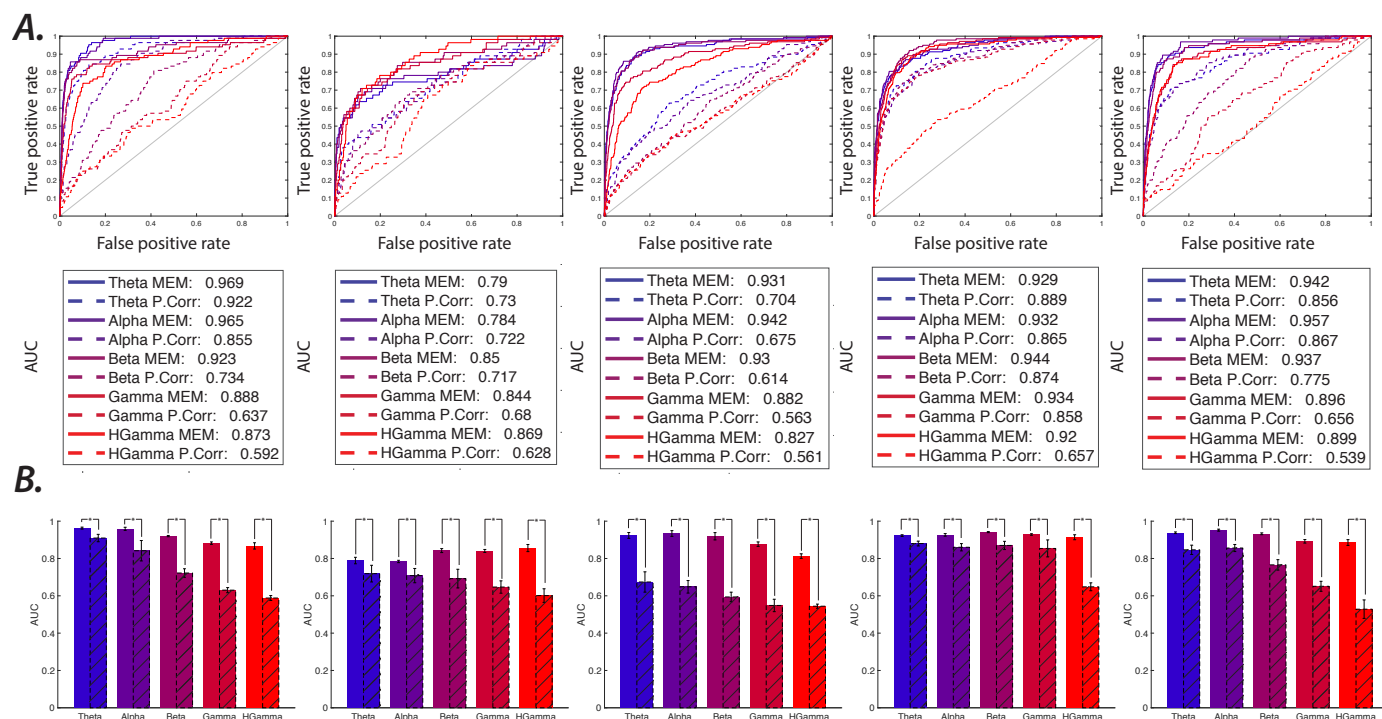


Figure 3. ROC analysis for the detection of SC based on functional interaction estimates. (A) ROC curves for the identification of anatomically connected versus unconnected ROI pairs (i.e., binarized SC) based on the average (14.6 ± 2.5 hours of one-hour segment data per patient) J_{ij} interaction matrices (solid lines) and average power FC (dashed lines) for all 5 patients (columns). Lines are color-coded based on frequency band. The bottom tables provide the AUC values for all curves. (B) The color-coded bars represent the mean AUC and error bars represent the standard deviation (across all one-hour segments) calculated for the SC matrices for different frequency bands. The average AUC calculated using J_{ij} interaction matrices (solid bars, left) are significantly (t -test, $p < 0.05$ FDR, marked by ‘*’) higher than that calculated using power FC matrices (dashed bars, right) across all frequency bands. Similar results are found at even higher SC pruning thresholds, as shown in SI-Figure 3.

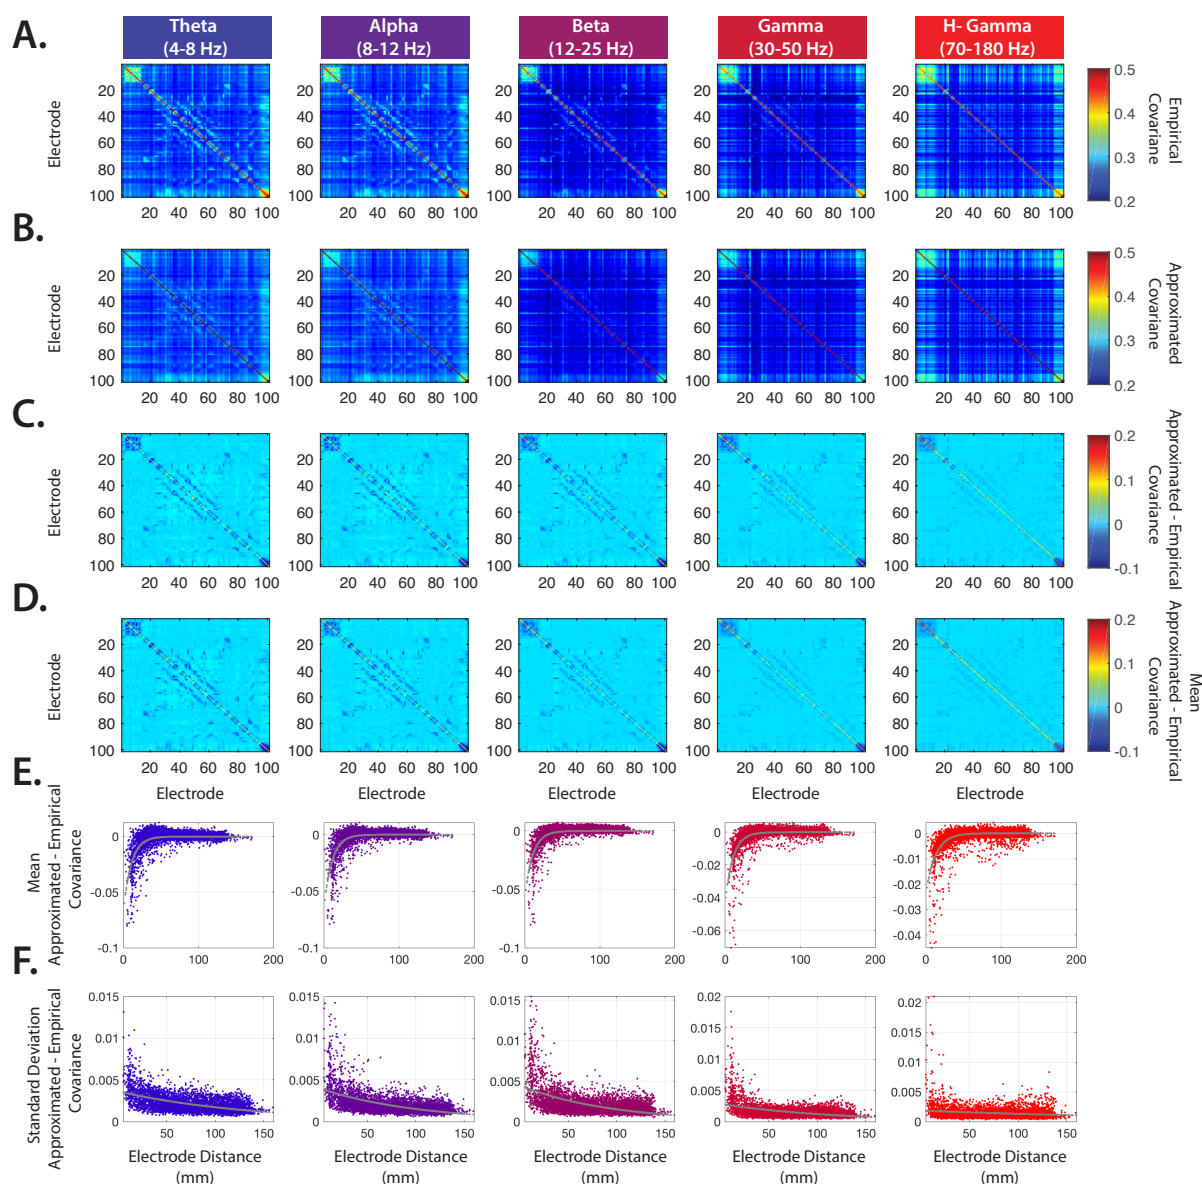


Figure 4. Comparison of covariance matrices of neural activity with those predicted by the pairwise MEM. (A) The covariance matrices calculated from band-passed power amplitude states from a single representative patient (#5). Each column contains results obtained from a distinct frequency band. (B) The covariance matrices estimated from the pairwise MEM. (C) Differences between the empirically observed covariance matrices and the computationally estimated covariance matrices obtained from the pairwise MEM for a single representative patient (#5). (D) The average difference matrices across all 19 data sets, reflecting the approximation error. (E) The average approximation error as a function of the pairwise Euclidean distance between electrodes, fit by an exponential function (green dashed line). Note that the error is notably higher between nearby electrodes. For convergent results across all subjects, see SI-Figure 9. (F) The standard deviation of the approximation error as a function of the pairwise Euclidean distance between electrodes.

Another way to assess the goodness of fit for the pairwise MEM is to calculate the degree to which the estimated probabilities for the observed activity states diverge from the empirically measured probabilities. As a baseline, this divergence between the pairwise model and the observed statistics is often compared to the divergence of the simpler first-order model, which assumes that each brain region is behaving independently from the others. It is important to note that this calculation can be computationally expensive for large systems because the state space increases exponentially ($= 2^n$) with the number of electrodes (n). To circumvent this issue, we instead compared the probability of all empirically observed states to their predicted probabilities. SI-Figure 10 demonstrates the degree of divergence between the empirical and estimated probabilities of all empirically observed states for a sample one-hour segment across all frequency bands. As expected, the largest estimation errors are observed in states with the lowest empirical probabilities. However the model produces notably large errors for a large number of states with relatively high empirical probabilities, specifically in the γ and high- γ bands (SI-Figure 10.A). Additionally, we examined the goodness of fit of the model for higher binarization thresholds of the power amplitude time series (SI-Figure 10.A). These results show that although higher thresholds provide different distributions of approximation errors, the model still provides a relatively poor prediction for a large number of states, consistent with the presence of higher-order interactions. It will be interesting in future to assess the presence and extent of such higher-order interaction in the γ and high- γ bands using higher-order MEMs [Ganmor et al. \(2011\)](#). Nevertheless, the ability of the estimated interaction matrices to predict the underlying SC is highly robust to all binarization thresholds, despite the difference in the patterns of approximation errors (SI-Figure 10.B-C and SI-Figure 11). Together, these results suggest that a simple model using only a combination of regional activation rates and pairwise co-activation rates can closely approximate the observed covariance structure in the iEEG power time-series.

Discussion

Here, we constructed a series of pairwise MEMs capable of reproducing the covariance structure observed in iEEG power amplitude states recorded from 5 patients with partial-onset refractory epilepsy. Careful examination of the covariance matrices predicted by the pairwise MEMs revealed markedly small approximation errors across all frequency bands. Nevertheless, we identified divergence between the estimated and empirical probabilities of states, mostly in the γ bands, suggesting that γ oscillations may be strongly driven by higher-order interactions. These results suggest that, while external common sources and higher-order interactions may contribute to the observed neural activity, the activation and pairwise co-activation rates between brain regions capture most of the observed variance. Importantly, our results also revealed that the inferred pairwise interactions in the MEM were highly stable over time and maintained a high degree of similarity across frequency bands. Notably, the estimated interaction matrices were both more consistent over time and more similar across frequencies than a commonly used measure of FC, the Pearson correlation between the power amplitude time series. We then investigated the relationship between the pairwise maximum entropy interactions, which were calculated only to account for the functional iEEG data, and the underlying white matter connectivity between regions by performing ROC analysis. AUC results revealed that the pairwise MEM provides accurate predictions of the underlying structural connectivity across all subjects and frequency bands, an effect similar to – though significantly stronger than – that observed in power amplitude FC at low frequencies. Taken together, these findings demonstrate the utility of the pairwise MEM for explaining covariance patterns in iEEG signals and in uncovering their underlying anatomical substrates.

The structural connectivity scaffold for functional connectivity

A critical open question in neuroscience is how the relatively fixed white matter structure of the brain gives rise to complex functional dynamics. Some have argued that these complex dynamics are well-approximated by a single underlying process ([Laumann et al., 2016](#); [Liegeois et al., 2017](#)), while others favor a model comprised of multiple processes that the brain can switch between ([Friston et al., 2003](#); [Deco et al., 2013](#); [Laumann et al., 2016](#)). In either scenario, it is critical to understand from a statistical perspective whether the observed dynamics can be explained by a fixed structural scaffold, such as that represented by the pattern of white matter fiber bundles connecting large-scale brain regions. The answer depends to some degree on whether or not one can identify statistical features of the functional dynamics that display relatively time-invariant properties. Some evidence for such time-invariance has recently been obtained from electrocorticography, where patterns of functional connectivity between electrodes appear relatively stable in windows larger than a few minutes ([Chu et al., 2012](#); [Kramer et al., 2010, 2011](#)). These data are consistent with our observations that patterns of interactions estimated over one-hour windows were relatively stable across several hours. Notably, this temporal stability was markedly higher in the interactions estimated from the pairwise MEM than in the functional connections estimated from the Pearson correlation in power amplitude time series. Our results therefore provide evidence of high stationarity of the intrinsic pairwise relationships between brain regions that gives rise to power amplitude dynamics, and further motivate a careful investigation into the correspondence between those relationships and the underlying anatomical projections.

A second line of evidence supporting the existence of a common structural scaffold for the observed dynamics is the

presence of notable similarity in the estimated pairwise interactions across frequency bands. This statistical similarity is reminiscent of the cross-frequency interactions that have been observed in other studies, and that are thought to play a role in integrating information across brain structures that are functionally specialized and spatially segregated (Canolty et al., 2006; Hyafil et al., 2015; Lisman and Jensen, 2013). It is commonly thought that the diverse rhythms present in different frequency bands are associated with different spatio-temporal scales of neural activity (Raghavachari et al., 2006; Canolty et al., 2007), with low frequencies driving activity over larger spatial areas and high frequencies driving activity over smaller spatial areas (Von Stein and Sarnthein, 2000). Our work offers a useful complement to these prior studies by demonstrating that the pairwise MEM provides higher cross-frequency similarity in estimated interactions than conventional measures of functional connectivity. Future work could examine the real-time dynamics of cross-frequency coupling by combining power amplitude states across all bands in a single model. Taken together, our observations of high cross-frequency similarity indicate that a common underlying physiological (Manning et al., 2009) or anatomical (Betz et al., 2018) mechanism is likely driving the functional interactions between brain regions across all frequency bands.

Our observation of the relatively high cross-frequency similarity and stationarity of interaction matrices estimated from the pairwise MEM is consistent with the notion that the patterns of interactions track an underlying structural scaffold. We validate this possibility with rigorous statistical testing, showing that the interactions can be used to predict underlying white matter structure in the same patients, and that these predictions are stronger for the pairwise MEM interactions than for a conventional measure of functional connectivity, consistent with prior work in other imaging modalities (Watanabe et al., 2013). The higher structure-function coupling obtained with the MEM could be due to the fact that the MEM simultaneously fits all regional activation rates and all pairwise interactions. Partial correlation approaches that operationalize the latter also tend to show high structure-function coupling (Salvador et al., 2005; Watanabe et al., 2013). It will be interesting in future work to consider the relevance of the MEM for explaining the effects of stimulation, for example where single external pulse stimulation of cortex propagates to mono- and poly-synaptically connected regions via cortico-cortical and cortico-subcortical projections (Steriade and Amzica, 1996; Matsumoto et al., 2006). In sum, our results suggest that the functional interactions facilitated by the direct underlying structural connectivity between brain regions largely shapes the complex broadband functional dynamics evident in power amplitude time series.

Statistical accuracy of the pairwise MEM and potential biophysical mechanisms

Pairwise MEMs have been employed effectively to explain patterns of collective neural activity across a range of spatial scales (Schneidman et al., 2006; Watanabe et al., 2013; Ashourvan et al., 2017). Here, we demonstrate that the pairwise MEM allow us to approximate the empirical covariance patterns of power amplitude states in several biologically relevant frequency bands. We observed that the accuracy of the approximated covariance in the pairwise MEM was markedly high across all frequency bands. While the overall accuracy was high, we sought to better understand where and why the pairwise MEM produced its greatest errors. We observed that the model systematically overestimated electrode activation rates while underestimating electrode-to-electrode co-activation rates, particularly between electrodes that were located in close physical proximity, the latter observation being consistent with prior work (Raghavachari et al., 2006; Canolty et al., 2007; Nunez and Srinivasan, 2006; Dickson et al., 2000). Additional examination of the empirical and estimated probabilities revealed notable divergence in the probabilities of a number of the observed states, with this divergence being most pronounced in the γ and high γ bands. Although some approximation error is expected due to our heuristic algorithm for calculating pairwise MEM parameters, we speculate that the error that we observe might also originate from biophysical mechanisms contributing to the co-activation of nearby electrodes. Examples of such biophysical mechanisms include the well-known and markedly strong white matter connectivity between nearby brain regions (Betz et al., 2018), spatially distributed activation sources, and spiral waves (Huang et al., 2010). Moreover, it is important to note that even though iEEG electrodes cover a large portion of the cortical surface in these patients, we are still only able to capture a small fraction of the full system, and it is highly likely that some amount of co-activation is driven by input from external sources that are not covered by an electrode. Although the MEM could be extended to include higher-order parameters beyond the pairwise interactions considered here, the practical limitation of partial brain coverage and the previously noted biophysical mechanisms combine to hinder the interpretability of higher-order mechanisms.

Methodological considerations

One common concern regarding the use of iEEG datasets from patients with epilepsy to study the functional organization of the human brain is their documented aberrations of structural connectivity as well as structure-function coupling (Liao et al., 2010; Besson et al., 2014; Bonilha et al., 2012; Zhang et al., 2011). Since iEEG datasets from healthy controls do not exist, we cannot directly address these concerns. Nevertheless, we draw on recent work demonstrating that the patterns of ECoG functional connectivity in patients show statistical similarities to structural connectivity estimated in healthy volunteers (Betz et al., 2018), and that this statistical similarity is upheld and even strengthened during ictal epochs (Reddy et al., 2018; Shah

et al., 2018). Future work using data from patients with other pathologies, or using source-localized MEG in healthy patients, could be helpful in further understanding the nature of the structure-function correspondence accessible to the MEM.

Preprocessing methods and parameter choices can impact estimates of structural connectivity. However, we demonstrate in SI-Figure 12 that the observed correspondence between functional interactions and structural connections were robust to reasonable variation in these choices. Although the pairwise MEM demonstrates a closer correspondence than the Pearson-based FC in predicting the SC across several different DTI processing methods, our results demonstrate that the deterministic streamline count matrix reflecting SC produces the strongest structure-function coupling. Future work could test the robustness of our findings to other important preprocessing parameters such as regional size, high-pass frequency filtering of the iEEG signals, and binarization threshold of power amplitude states.

It is of interest to determine the contribution of distance to structure-function coupling. The strength of a structural connection between two brain regions tends to be negatively correlated with the distance between them, as does the strength of the functional connection between them. One possibility is that structure-function coupling is driven by a common source that is simultaneously measured by nearby electrodes. To discount this possibility at least partially, we demonstrate that the pairwise MEM can be used to accurately identify structural connections between distant electrode pairs. This observation suggests that the functional dynamics captured by the pairwise MEM extends beyond common sources or local spreading phenomena, and that the presence of a structural connection between brain regions likely plays a crucial role in shaping the emergent functional activity across a wide range of frequencies. Nonetheless, future work should aim to validate our observations and further tease out the effect of electrode distance using datasets from patients with more uniform electrode coverage.

Conclusion

Our observations have potentially important implications for understanding large-scale functional brain dynamics as well as our ability to modulate these dynamics via stimulation or resection. We demonstrated that defining the band-passed power of the iEEG signal as a functional brain state allowed us to uncover strong structure-function coupling in large-scale networks. Therefore our results provide converging evidence of the functional relevance of iEEG signal power. Furthermore our results demonstrate that the apparent divergence between structural and functional connectivity is minimized when the intrinsic functional relationships between brain regions are not examined as independent pairs. These findings are consistent with the notion that the pairwise MEM may be particularly sensitive to structurally-driven functional relations while conventional functional connectivity methods may be particularly sensitive to non-structurally-driven functional relations that might vary appreciably over short time intervals. Finally, the observed high degree of structure-function coupling suggests that structural connectivity is a useful proxy for time-invariant functional relationships (Betzel et al., 2018; Reddy et al., 2018). This observation could be useful in the treatment of epilepsy patients, where access to the brain is traditionally limited to recording loci but could be augmented with non-invasive measurements of structural connectivity for more informed surgical planning (Shah et al., 2018). Indeed, it is intuitively plausible that computational models built to inform the modulation of abnormal functional dynamics via stimulation (Gu et al., 2015) or resection (Khambhati et al., 2016) may be able to utilize patient specific structural connectivity in place of or to augment patient specific functional connectivity.

Materials and Methods

Patient information

5 patients (mean age 41.6, standard deviation 4.8; 3 female) undergoing surgical treatment for medically refractory epilepsy underwent implantation of subdural electrodes for localization of the seizure onset zone. All patients had unilateral temporal lobe epilepsy, determined by comprehensive clinical evaluation and validated by seizure free one-year outcomes following temporal lobectomy.

Intracranial EEG Acquisition

De-identified patient data was downloaded from the online International Epilepsy Electrophysiology Portal (iEEG Portal, <http://www.ieeg.org>) (Wagenaar et al., 2013). iEEG signals were recorded at 512 Hz at the Hospital of the University of Pennsylvania, Philadelphia, PA. Subdural electrode (Ad Tech Medical Instruments, Racine, WI) configurations consisted of linear strip and two-dimensional grid arrays (2.3 mm diameter with 10 mm inter-contact spacing). A referential montage with reference electrode distant to the seizure onset zone was used for recording iEEG signals. Inter-ictal periods were selected to be at least 6 hours away from clinically marked seizures.

Image Acquisition

MRI data were acquired using a 3-Tesla Siemens scanner (Siemens Magnetom Trio Tim Syngo MR B17, Germany) at the Hospital of the University of Pennsylvania. Diffusion-weighted images were acquired in a single-shot echo-planar imaging,

multi-shell protocol ($2.5 \times 2.5 \times 2.5 \text{ mm}^3$ resolution, TR = 5216 ms, TE = 100 ms, FOV = 220 mm, MB acceleration factor = 2, flip/refocus angle = 90/180 degrees, phase encoding direction = anterior to posterior). A total of 119 volumes were acquired at $b=0$ (16 volumes), $b=300$ (8 volumes), $b=700$ (31 volumes), and $b=2000$ (64 volumes). T1-weighted MPRAGE images were also obtained for each subject ($0.94 \times 0.94 \times 1.0 \text{ mm}^3$ resolution, TR = 1810 ms, TE = 3.51 ms, FOV = 240 mm, flip angle = 9 degrees, phase encoding direction = right to left).

Structural connectivity

T1-weighted MPRAGE images were used to co-register MNI-space atlases to subject structural space via AntsRegistrationSyN (Avants et al., 2008). Similarly, we used the first b0 image from each patient's diffusion sequence to calculate co-registration transforms from subject-specific diffusion to structural space. Using this set of transforms, we brought individualized AAL 600 atlases (Hermundstad et al., 2013) for all five patients into diffusion space for tractography. Atlas to diffusion co-registrations were performed using ANTs and FSLs FLIRT. Eddy current and motion corrections were performed using FSLs 5.0.9 EDDY patch. To mitigate susceptibility distortions, we used the subjects MPRAGE T1 structural scan. This image was first brought to diffusion space by using the inverse of our FLIRT transformation for each subject, and then contrast-inverted and intensity-matched to the DWI image using FSLMATHS. Finally, the DWI image was non-linearly transformed to the shape of the MPRAGE scan, as these acquisitions are not subject to the same susceptibility distortions observed in DWIs.

We evaluated several distinct methodological approaches, and we also assessed the consistency of our findings across them. We performed deterministic and probabilistic tractography in Camino (Cook et al., 2006) on distortion-corrected DWIs to construct adjacency matrices for structural connectivity. Deterministic tractography parameters were stringent so we could err on the side of caution in evaluating the physiological feasibility of our reconstructed white matter connections. These parameters included removing estimated fibers in areas of the brain where 90% or greater of diffusion was estimated to be isotropic versus restricted (Zhan et al., 2015), where fractional anisotropy was determined to be below 0.05, or where a T1-derived brain mask ended. We also removed fibers that curved more than 50 degrees over 5 mm to again restrict fiber reconstructions to physiological feasibility. Finally, we removed fibers shorter than 10 mm to minimize spurious short-range connections based on only a few voxels, and we also removed fibers over 100 mm to prevent looping artifactual fibers from impacting our results. Structural adjacency matrices were constructed from the number of streamlines that begin and end in each pre-defined ROI. We employed equivalent parameters for probabilistic tractography and utilized the PICO algorithm (Parker et al., 2003). Processing power limited us to seeding 400 times per voxel. Within deterministic and probabilistic tractography, we evaluated streamline count (SC), fractional anisotropy (FA), neurite density (ND), and orientation dispersion (OD) weighted adjacency matrices.

Pairwise maximum entropy model

To form unbiased predictions for the probabilities of various functional brain states, we fit a pairwise maximum entropy model, which is motivated by the principle of maximum entropy. The principle of maximum entropy states that when estimating a probability distribution given some desired constraints, one ought to consider the distribution that maximizes the uncertainty (i.e., entropy); choosing any other distribution that lowers the entropy would assume additional information beyond our desired constraints. Fitting the pairwise MEM entails tuning the first and second-order interaction parameters between regions so that the predicted activation rates and co-activation rates match the empirically observed values. An accurate fit of the pairwise MEM implies that the observed patterns of collective activity can be understood as emerging from each region's independent activation rate combined with regions' joint activation rates. In other words, the pairwise MEM allows us to establish a model of iEEG functional power dynamics as a probabilistic process shaped by underlying pairwise relationships between brain regions.

We fit the pairwise MEM to the thresholded and binarized normalized power amplitude states of electrodes during one hour intervals of recording per subject (Watanabe et al., 2013). To calculate the power amplitude states, we first segmented the time-series into 36, 100-second windows (SI-Figure 1.A). Next, for each segment and each electrode we calculated the wavelet power (Grinsted et al., 2004) of the time-series (Figure 1.B). We band-passed the wavelet power by averaging the power amplitude across all the frequencies within each band. Next, the power amplitude time-series (all bands) were high-pass filtered at 0.5 Hz to minimize the effect of ultra slow fluctuations of power on the binarized states. The high-passed time series were then z-scored and binarized by thresholding at zero to create band-passed power amplitude states (Figure 1.C).

At each time t , the power amplitude state is defined by $V^t = [\sigma_1^t, \sigma_2^t, \dots, \sigma_{i-1}^t, \sigma_{i+1}^t, \dots, \sigma_N^t]$, where σ_i^t is the binarized power amplitude of electrode i at time t , where $\sigma_i^t = 1$ (0) for power above (below) the threshold, and N is the total number of electrodes. For electrode i , the empirical power amplitude activation rate $\langle \sigma_i \rangle$ is given by $(\frac{1}{T}) \sum_{t=1}^T \sigma_i^t$, where T is the number of time slices. Likewise, the empirical co-activation rate between electrodes i and j , $\langle \sigma_i \sigma_j \rangle$, is defined by $(\frac{1}{T}) \sum_{t=1}^T \sigma_i^t \sigma_j^t$.

Here, our only constraints were that the model averages $\langle \sigma_i \rangle_m$ and $\langle \sigma_i \sigma_j \rangle_m$ matched the empirical values of $\langle \sigma_i \rangle$ and $\langle \sigma_i \sigma_j \rangle$, respectively. Given these constraints, it is known that the probability distribution that maximizes the entropy is the

Boltzmann distribution (Jaynes, 1957):

$$P(V_k) = e^{-E(V_k)} / \sum_{q=1}^{2^N} e^{-E(V_q)}, \quad (1)$$

where $P(V_k)$ is the probability distribution of the k -th state V_k , and $E(V_k)$ is the energy of this state, which is given by:

$$E(V_k) = - \sum_{i=1}^N h_i \sigma_i(V_k) - \frac{1}{2} \sum_{i,j=1}^N J_{ij} \sigma_i(V_k) \sigma_j(V_k), \quad (2)$$

where $\sigma_i(V_k)$ is the value of σ_i for state V_k , h_i represents the expected base power amplitude rate of electrode i in isolation, and J_{ij} represents the functional interaction between electrodes i and j .

Fitting the pairwise MEM entails iterative adjustment of the parameters h_i and J_{ij} with a gradient descent algorithm (Watanabe et al., 2014) until the empirical averages $\langle \sigma_i \rangle$ and $\langle \sigma_i \sigma_j \rangle$ match those in the model, namely $\langle \sigma_i \rangle_m = \sum_{q=1}^{2^{n-1}} \sigma_i(V_q) P(V_q)$ and $\langle \sigma_i \sigma_j \rangle_m = \sum_{q=1}^{2^{n-1}} \sigma_i(V_q) \sigma_j(V_q) P(V_q)$. Since the state space is prohibitively large in our data, we followed (Gu et al., 2018) by approximating the model averages $\langle \sigma_i \rangle_m$ and $\langle \sigma_i \sigma_j \rangle_m$ by calculating the covariance of a sequence of random samples ($N = 10300000$, after discarding the first 300000 and downsampling by a factor of 500) from the probability distribution using the Metropolis–Hastings algorithm (Metropolis et al., 1953).

Comparison between structural connectivity and functional connectivity

We conducted ROC analysis to determine the accuracy with which the strength of the functional interactions in the pairwise MEM predict the presence of structural connectivity. We constructed the ROC curves to represent the true positive (structurally connected brain regions with functional interaction weight above the threshold) and false positive (structurally disconnected brain regions with functional interaction weight above the threshold) at different detection thresholds. The area under the curve (AUC) of the ROC plot represents the accuracy of the classification, where an AUC value of 1 indicates perfect classification, and an AUC value of 0.5 indicates classification performance at the chance level.

Acknowledgements

This work was supported by an NINDS 1R01NS099348 to Litt and Bassett (MPI). BL also acknowledges support from the Mirowski Foundation and from Neil Barbara Smit. DSB also acknowledges support from the John D. and Catherine T. MacArthur Foundation, the Alfred P. Sloan Foundation, the Army Research Office through contract number W911NF-14-1-0679, the Army Research Laboratory through contract number W911NF-10-2-0022, the National Institute of Health (2-R01-DC-009209-11, 1R01HD086888-01, R01-MH107235, R01-MH107703, R01MH109520, 1R01NS099348 and R21-MH-106799), the Office of Naval Research, and the National Science Foundation (BCS-1441502, CAREER PHY-1554488, BCS-1631550, and CNS-1626008). The content is solely the responsibility of the authors and does not necessarily represent the official views of any of the funding agencies.

References

- Yusuke Adachi, Takahiro Osada, Olaf Sporns, Takamitsu Watanabe, Teppei Matsui, Kentaro Miyamoto, and Yasushi Miyashita. Functional connectivity between anatomically unconnected areas is shaped by collective network-level effects in the macaque cortex. *Cerebral cortex*, 22(7):1586–1592, 2011.
- Arian Ashourvan, Shi Gu, Marcelo G Mattar, Jean M Vettel, and Danielle S Bassett. The energy landscape underpinning module dynamics in the human brain connectome. *Neuroimage*, 157:364–380, 2017.
- Brian B Avants, Charles L Epstein, Murray Grossman, and James C Gee. Symmetric diffeomorphic image registration with cross-correlation: evaluating automated labeling of elderly and neurodegenerative brain. *Medical image analysis*, 12(1): 26–41, 2008.
- D S Bassett and E Bullmore. Small-world brain networks. *Neuroscientist*, 12(6):512–523, 2006.
- D S Bassett and O Sporns. Network neuroscience. *Nat Neurosci*, 20(3):353–364, 2017.
- D S Bassett, P Zurn, and J I Gold. On the nature and use of models in network neuroscience. *Nat Rev Neurosci*, Epub ahead of print, 2018.

- Danielle S Bassett, Edward Bullmore, Beth A Verchinski, Venkata S Mattay, Daniel R Weinberger, and Andreas Meyer-Lindenberg. Hierarchical organization of human cortical networks in health and schizophrenia. *Journal of Neuroscience*, 28(37):9239–9248, 2008.
- Danielle S Bassett, Nicholas F Wymbs, Mason A Porter, Peter J Mucha, Jean M Carlson, and Scott T Grafton. Dynamic reconfiguration of human brain networks during learning. *Proceedings of the National Academy of Sciences*, 108(18):7641–7646, 2011.
- Danielle S Bassett, Muzhi Yang, Nicholas F Wymbs, and Scott T Grafton. Learning-induced autonomy of sensorimotor systems. *Nature neuroscience*, 18(5):744, 2015.
- Cassiano O Becker, Sérgio Pequito, George J Pappas, Michael B Miller, Scott T Grafton, Danielle S Bassett, and Victor M Preciado. Spectral mapping of brain functional connectivity from diffusion imaging. *Scientific reports*, 8(1):1411, 2018.
- Pierre Besson, Vera Dinkelacker, Romain Valabregue, Lionel Thivard, Xavier Leclerc, Michel Baulac, Daniela Sammler, Olivier Colliot, Stéphane Lehericy, Séverine Samson, et al. Structural connectivity differences in left and right temporal lobe epilepsy. *Neuroimage*, 100:135–144, 2014.
- R F Betzel and D S Bassett. Specificity and robustness of long-distance connections in weighted, interareal connectomes. *Proc Natl Acad Sci U S A*, 115(21):E4880–E4889, 2018.
- R F Betzel, J D Medaglia, A E Kahn, J Soffer, D R Schonhaut, and D S Bassett. Inter-regional ECoG correlations predicted by communication dynamics, geometry, and correlated gene expression. *arXiv*, 1706:06088, 2018.
- Leonardo Bonilha, Travis Nesland, Gabriel U Martz, Jane E Joseph, Maria V Spampinato, Jonathan C Edwards, and Ali Tabesh. Medial temporal lobe epilepsy is associated with neuronal fibre loss and paradoxical increase in structural connectivity of limbic structures. *J Neurol Neurosurg Psychiatry*, pages jnnp–2012, 2012.
- Urs Braun, Axel Schäfer, Henrik Walter, Susanne Erk, Nina Romanczuk-Seiferth, Leila Haddad, Janina I Schweiger, Oliver Grimm, Andreas Heinz, Heike Tost, et al. Dynamic reconfiguration of frontal brain networks during executive cognition in humans. *Proceedings of the National Academy of Sciences*, 112(37):11678–11683, 2015.
- Ryan T Canolty, Erik Edwards, Sarang S Dalal, Maryam Soltani, Srikantan S Nagarajan, Heidi E Kirsch, Mitchel S Berger, Nicholas M Barbaro, and Robert T Knight. High gamma power is phase-locked to theta oscillations in human neocortex. *science*, 313(5793):1626–1628, 2006.
- Ryan T Canolty, Maryam Soltani, Sarang S Dalal, Erik Edwards, Nina F Dronkers, Srikantan S Nagarajan, Heidi E Kirsch, Nicholas M Barbaro, and Robert T Knight. Spatiotemporal dynamics of word processing in the human brain. *Frontiers in neuroscience*, 1:14, 2007.
- Catherine J Chu, Mark A Kramer, Jay Pathmanathan, Matt T Bianchi, M Brandon Westover, Lauren Wizon, and Sydney S Cash. Emergence of stable functional networks in long-term human electroencephalography. *Journal of Neuroscience*, 32(8):2703–2713, 2012.
- Luis Concha, Donald W Gross, B Matt Wheatley, and Christian Beaulieu. Diffusion tensor imaging of time-dependent axonal and myelin degradation after corpus callosotomy in epilepsy patients. *NeuroImage*, 32(3):1090–1099, 2006. ISSN 10538119. doi: 10.1016/j.neuroimage.2006.04.187.
- Luis Concha, Christian Beaulieu, B. Matt Wheatley, and Donald W. Gross. Bilateral White Matter Diffusion Changes Persist after Epilepsy Surgery. *Epilepsia*, 48(5):931–940, may 2007. ISSN 0013-9580. doi: 10.1111/j.1528-1167.2007.01006.x.
- PA Cook, Y Bai, SKKS Nedjati-Gilani, KK Seunarine, MG Hall, GJ Parker, and DC Alexander. Camino: open-source diffusion-mri reconstruction and processing. In *14th scientific meeting of the international society for magnetic resonance in medicine*, volume 2759, page 2759. Seattle WA, USA, 2006.
- Gustavo Deco, Adrián Ponce-Alvarez, Dante Mantini, Gian Luca Romani, Patric Hagmann, and Maurizio Corbetta. Resting-state functional connectivity emerges from structurally and dynamically shaped slow linear fluctuations. *Journal of Neuroscience*, 33(27):11239–11252, 2013.
- Clayton T Dickson, Gerardo Biella, and Marco De Curtis. Evidence for spatial modules mediated by temporal synchronization of carbachol-induced gamma rhythm in medial entorhinal cortex. *Journal of Neuroscience*, 20(20):7846–7854, 2000.

- Chad J Donahue, Stamatios N Sotiropoulos, Saad Jbabdi, Moises Hernandez-Fernandez, Timothy E Behrens, Tim B Dyrby, Timothy Coalson, Henry Kennedy, Kenneth Knoblauch, David C Van Essen, et al. Using diffusion tractography to predict cortical connection strength and distance: a quantitative comparison with tracers in the monkey. *Journal of Neuroscience*, 36(25):6758–6770, 2016.
- Nico UF Dosenbach, Binyam Nardos, Alexander L Cohen, Damien A Fair, Jonathan D Power, Jessica A Church, Steven M Nelson, Gagan S Wig, Alecia C Vogel, Christina N Lessov-Schlaggar, et al. Prediction of individual brain maturity using fmri. *Science*, 329(5997):1358–1361, 2010.
- Pascal Fries. Neuronal gamma-band synchronization as a fundamental process in cortical computation. *Annual review of neuroscience*, 32:209–224, 2009.
- K.J. Friston, L. Harrison, and W. Penny. Dynamic causal modelling. *NeuroImage*, 19(4):1273 – 1302, 2003. ISSN 1053-8119.
- Elad Ganmor, Ronen Segev, and Elad Schneidman. Sparse low-order interaction network underlies a highly correlated and learnable neural population code. *Proceedings of the National Academy of Sciences*, 108(23):9679–9684, jun 2011. ISSN 0027-8424. doi: 10.1073/pnas.1019641108.
- Joaquín Goñi, Martijn P van den Heuvel, Andrea Avena-Koenigsberger, Nieves Velez de Mendizabal, Richard F Betzel, Alessandra Griffa, Patric Hagmann, Bernat Corominas-Murtra, Jean-Philippe Thiran, and Olaf Sporns. Resting-brain functional connectivity predicted by analytic measures of network communication. *Proceedings of the National Academy of Sciences*, 111(2):833–838, 2014.
- Roberta Grech, Tracey Cassar, Joseph Muscat, Kenneth P Camilleri, Simon G Fabri, Michalis Zervakis, Petros Xanthopoulos, Vangelis Sakkalis, and Bart Vanrumste. Review on solving the inverse problem in eeg source analysis. *Journal of neuroengineering and rehabilitation*, 5(1):25, 2008.
- A Grinsted, J C Moore, and S Jevrejeva. Application of the cross wavelet transform and wavelet coherence to geophysical time series. *Nonlin. Processes Geophys.*, 11:561–566, 2004.
- Shi Gu, Fabio Pasqualetti, Matthew Cieslak, Qawi K Telesford, B Yu Alfred, Ari E Kahn, John D Medaglia, Jean M Vettel, Michael B Miller, Scott T Grafton, et al. Controllability of structural brain networks. *Nature communications*, 6:8414, 2015.
- Shi Gu, Matthew Cieslak, Benjamin Baird, Sarah F Muldoon, Scott T Grafton, Fabio Pasqualetti, and Danielle S Bassett. The energy landscape of neurophysiological activity implicit in brain network structure. *Scientific reports*, 8(1):2507, 2018.
- Patric Hagmann, Leila Cammoun, Xavier Gigandet, Reto Meuli, Christopher J Honey, Van J Wedeen, and Olaf Sporns. Mapping the structural core of human cerebral cortex. *PLoS biology*, 6(7):e159, 2008.
- Patric Hagmann, Olaf Sporns, Neel Madan, Leila Cammoun, Rudolph Pienaar, Van Jay Wedeen, Reto Meuli, J-P Thiran, and PE Grant. White matter maturation reshapes structural connectivity in the late developing human brain. *Proceedings of the National Academy of Sciences*, 107(44):19067–19072, 2010.
- A M Hermundstad, K S Brown, D S Bassett, E M Aminoff, A Frithsen, A Johnson, C M Tipper, M B Miller, S T Grafton, and J M Carlson. Structurally-constrained relationships between cognitive states in the human brain. *PLoS Comput Biol*, 10(5): e1003591, 2014.
- Ann M Hermundstad, Danielle S Bassett, Kevin S Brown, Elissa M Aminoff, David Clewett, Scott Freeman, Amy Frithsen, Arianne Johnson, Christine M Tipper, Michael B Miller, et al. Structural foundations of resting-state and task-based functional connectivity in the human brain. *Proceedings of the National Academy of Sciences*, 110(15):6169–6174, 2013.
- CJ Honey, O Sporns, Leila Cammoun, Xavier Gigandet, Jean-Philippe Thiran, Reto Meuli, and Patric Hagmann. Predicting human resting-state functional connectivity from structural connectivity. *Proceedings of the National Academy of Sciences*, 106(6):2035–2040, 2009.
- Xiaoying Huang, Weifeng Xu, Jianmin Liang, Kentaroh Takagaki, Xin Gao, and Jian-young Wu. Spiral wave dynamics in neocortex. *Neuron*, 68(5):978–990, 2010.
- Alexandre Hyafil, Anne-Lise Giraud, Lorenzo Fontolan, and Boris Gutkin. Neural cross-frequency coupling: connecting architectures, mechanisms, and functions. *Trends in neurosciences*, 38(11):725–740, 2015.

- Corentin Jacques, Nathan Witthoft, Kevin S Weiner, Brett L Foster, Vinitha Rangarajan, Dora Hermes, Kai J Miller, Josef Parvizi, and Kalanit Grill-Spector. Corresponding ecog and fmri category-selective signals in human ventral temporal cortex. *Neuropsychologia*, 83:14–28, 2016.
- Edwin T Jaynes. Information theory and statistical mechanics. *Physical review*, 106(4):620, 1957.
- Marcus Kaiser and Claus C Hilgetag. Modelling the development of cortical systems networks. *Neurocomputing*, 58:297–302, 2004.
- Ankit N Khambhati, Kathryn A Davis, Timothy H Lucas, Brian Litt, and Danielle S Bassett. Virtual cortical resection reveals push-pull network control preceding seizure evolution. *Neuron*, 91(5):1170–1182, 2016.
- Mark A Kramer, Uri T Eden, Eric D Kolaczyk, Rodrigo Zepeda, Emad N Eskandar, and Sydney S Cash. Coalescence and fragmentation of cortical networks during focal seizures. *Journal of Neuroscience*, 30(30):10076–10085, 2010.
- Mark A Kramer, Uri T Eden, Kyle Q Lepage, Eric D Kolaczyk, Matt T Bianchi, and Sydney S Cash. Emergence of persistent networks in long-term intracranial eeg recordings. *Journal of Neuroscience*, 31(44):15757–15767, 2011.
- Jean-Philippe Lachaux, Pierre Fonlupt, Philippe Kahane, Lorella Minotti, Dominique Hoffmann, Olivier Bertrand, and Monica Baciú. Relationship between task-related gamma oscillations and bold signal: New insights from combined fmri and intracranial eeg. *Human brain mapping*, 28(12):1368–1375, 2007.
- Timothy O Laumann, Abraham Z Snyder, Anish Mitra, Evan M Gordon, Caterina Gratton, Babatunde Adeyemo, Adrian W Gilmore, Steven M Nelson, Jeff J Berg, Deanna J Greene, et al. On the stability of bold fmri correlations. *Cerebral cortex*, 27(10):4719–4732, 2016.
- John D Lewis, Rebecca J Theilmann, Martin I Sereno, and Jeanne Townsend. The relation between connection length and degree of connectivity in young adults: a dti analysis. *Cerebral Cortex*, 19(3):554–562, 2008.
- Wei Liao, Zhiqiang Zhang, Zhengyong Pan, Dante Mantini, Jurong Ding, Xujun Duan, Cheng Luo, Guangming Lu, and Huaifu Chen. Altered functional connectivity and small-world in mesial temporal lobe epilepsy. *PloS one*, 5(1):e8525, 2010.
- Raphael Liegeois, Timothy O Laumann, Abraham Z Snyder, Juan Zhou, and BT Thomas Yeo. Interpreting temporal fluctuations in resting-state functional connectivity mri. *Neuroimage*, 2017.
- John E Lisman and Ole Jensen. The theta-gamma neural code. *Neuron*, 77(6):1002–1016, 2013.
- Nikos K Logothetis, Jon Pauls, Mark Augath, Torsten Trinath, and Axel Oeltermann. Neurophysiological investigation of the basis of the fmri signal. *Nature*, 412(6843):150, 2001.
- Jie Lu, Hesheng Liu, Miao Zhang, Danhong Wang, Yanxiang Cao, Qingfeng Ma, Dongdong Rong, Xiaoyi Wang, Randy L Buckner, and Kuncheng Li. Focal pontine lesions provide evidence that intrinsic functional connectivity reflects polysynaptic anatomical pathways. *Journal of Neuroscience*, 31(42):15065–15071, 2011.
- Jeremy R Manning, Joshua Jacobs, Itzhak Fried, and Michael J Kahana. Broadband shifts in local field potential power spectra are correlated with single-neuron spiking in humans. *Journal of Neuroscience*, 29(43):13613–13620, 2009.
- Riki Matsumoto, Dileep R Nair, Eric LaPresto, William Bingaman, Hiroshi Shibasaki, and Hans O Lüders. Functional connectivity in human cortical motor system: a cortico-cortical evoked potential study. *Brain*, 130(1):181–197, 2006.
- Vinod Menon. Large-scale brain networks and psychopathology: a unifying triple network model. *Trends in cognitive sciences*, 15(10):483–506, 2011.
- Nicholas Metropolis, Arianna W Rosenbluth, Marshall N Rosenbluth, Augusta H Teller, and Edward Teller. Equation of state calculations by fast computing machines. *The journal of chemical physics*, 21(6):1087–1092, 1953.
- David Meunier, Sophie Achard, Alexa Morcom, and Ed Bullmore. Age-related changes in modular organization of human brain functional networks. *Neuroimage*, 44(3):715–723, 2009a.
- David Meunier, Renaud Lambiotte, Alex Fornito, Karen Ersche, and Edward T Bullmore. Hierarchical modularity in human brain functional networks. *Frontiers in neuroinformatics*, 3:37, 2009b.

- David Meunier, Renaud Lambiotte, and Edward T Bullmore. Modular and hierarchically modular organization of brain networks. *Frontiers in neuroscience*, 4:200, 2010.
- Sifis Micheloyannis, Ellie Pachou, Cornelis Jan Stam, Michael Breakspear, Panagiotis Bitsios, Michael Vourkas, Sophia Erimaki, and Michael Zervakis. Small-world networks and disturbed functional connectivity in schizophrenia. *Schizophrenia research*, 87(1):60–66, 2006.
- Yuval Nir, Roy Mukamel, Ilan Dinstein, Eran Privman, Michal Harel, Lior Fisch, Hagar Gelbard-Sagiv, Svetlana Kipervasser, Fani Andelman, Miri Y Neufeld, et al. Interhemispheric correlations of slow spontaneous neuronal fluctuations revealed in human sensory cortex. *Nature neuroscience*, 11(9):1100, 2008.
- Paul L Nunez and Ramesh Srinivasan. *Electric fields of the brain: the neurophysics of EEG*. Oxford University Press, USA, 2006.
- George A Ojemann, Nick F Ramsey, and Jeffrey Ojemann. Relation between functional magnetic resonance imaging (fmri) and single neuron, local field potential (lfp) and electrocorticography (ecog) activity in human cortex. *Frontiers in human neuroscience*, 7:34, 2013.
- Geoffrey JM Parker, Hamied A Haroon, and Claudia AM Wheeler-Kingshott. A framework for a streamline-based probabilistic index of connectivity (pico) using a structural interpretation of mri diffusion measurements. *Journal of Magnetic Resonance Imaging*, 18(2):242–254, 2003.
- Sridhar Raghavachari, John E Lisman, Michele Tully, Joseph R Madsen, EB Bromfield, and Michael J Kahana. Theta oscillations in human cortex during a working-memory task: evidence for local generators. *Journal of Neurophysiology*, 95(3):1630–1638, 2006.
- P G Reddy, R F Betzel, A N Khambhati, P Shah, L Kini, B Litt, T H Lucas, K A Davis, and D S Bassett. Genetic and neuroanatomical support for functional brain network dynamics in epilepsy. *arXiv*, 1809:03934, 2018.
- Yasser Roudi, Sheila Nirenberg, and Peter E Latham. Pairwise maximum entropy models for studying large biological systems: when they can work and when they can't. *PLoS Comput Biol*, 5(5):e1000380, 2009.
- Mikhail Rubinov, Rolf JF Ypma, Charles Watson, and Edward T Bullmore. Wiring cost and topological participation of the mouse brain connectome. *Proceedings of the National Academy of Sciences*, 112(32):10032–10037, 2015.
- Raymond Salvador, John Suckling, Martin R Coleman, John D Pickard, David Menon, and ED Bullmore. Neurophysiological architecture of functional magnetic resonance images of human brain. *Cerebral cortex*, 15(9):1332–1342, 2005.
- Elad Schneidman, Michael J Berry, Ronen Segev, and William Bialek. Weak pairwise correlations imply strongly correlated network states in a neural population. *Nature*, 440(7087):1007–1012, 2006.
- P Shah, A Ashourvan, F Mikhail, A Pines, L Kini, R T Shinohara, D S Bassett, B Litt, and K A Davis. Local structural connectivity directs seizure spread in focal epilepsy. *bioRxiv*, 406793, 2018.
- James M Shine, Patrick G Bissett, Peter T Bell, Oluwasanmi Koyejo, Joshua H Balsters, Krzysztof J Gorgolewski, Craig A Moodie, and Russell A Poldrack. The dynamics of functional brain networks: integrated network states during cognitive task performance. *Neuron*, 92(2):544–554, 2016.
- Jonathon Shlens, Greg D Field, Jeffrey L Gauthier, Matthew I Grivich, Dumitru Petrusca, Alexander Sher, Alan M Litke, and EJ Chichilnisky. The structure of multi-neuron firing patterns in primate retina. *The Journal of neuroscience*, 26(32):8254–8266, 2006.
- Pawel Skudlarski, Kanchana Jagannathan, Vince D Calhoun, Michelle Hampson, Beata A Skudlarska, and Godfrey Pearlson. Measuring brain connectivity: diffusion tensor imaging validates resting state temporal correlations. *Neuroimage*, 43(3):554–561, 2008.
- Pawel Skudlarski, Kanchana Jagannathan, Karen Anderson, Michael C Stevens, Vince D Calhoun, Beata A Skudlarska, and Godfrey Pearlson. Brain connectivity is not only lower but different in schizophrenia: a combined anatomical and functional approach. *Biological psychiatry*, 68(1):61–69, 2010.

- Sheng-Kwei Song, Shu-Wei Sun, Won-Kyu Ju, Shiow-Jiuan Lin, Anne H Cross, and Arthur H Neufeld. Diffusion tensor imaging detects and differentiates axon and myelin degeneration in mouse optic nerve after retinal ischemia. *NeuroImage*, 20(3):1714–1722, nov 2003. ISSN 1053-8119. doi: 10.1016/J.NEUROIMAGE.2003.07.005.
- Mircea Steriade and Florin Amzica. Intracortical and corticothalamic coherency of fast spontaneous oscillations. *Proceedings of the National Academy of Sciences*, 93(6):2533–2538, 1996.
- Kaustubh Supekar, Mark Musen, and Vinod Menon. Development of large-scale functional brain networks in children. *PLoS biology*, 7(7):e1000157, 2009.
- Peter N. Taylor, Nishant Sinha, Yujiang Wang, Sjoerd B. Vos, Jane de Tisi, Anna Miserocchi, Andrew W. McEvoy, Gavin P. Winston, and John S. Duncan. The impact of epilepsy surgery on the structural connectome and its relation to outcome. *NeuroImage: Clinical*, 18:202–214, jan 2018. ISSN 2213-1582. doi: 10.1016/J.NICL.2018.01.028.
- Peter J Uhlhaas and Wolf Singer. Neural synchrony in brain disorders: relevance for cognitive dysfunctions and pathophysiology. *Neuron*, 52(1):155–168, 2006.
- Martijn P Van den Heuvel and Olaf Sporns. An anatomical substrate for integration among functional networks in human cortex. *Journal of Neuroscience*, 33(36):14489–14500, 2013.
- Astrid Von Stein and Johannes Sarnthein. Different frequencies for different scales of cortical integration: from local gamma to long range alpha/theta synchronization. *International journal of psychophysiology*, 38(3):301–313, 2000.
- Joost B Wagenaar, Benjamin H Brinkmann, Zachary Ives, Gregory A Worrell, and Brian Litt. A multimodal platform for cloud-based collaborative research. In *Neural Engineering (NER), 2013 6th International IEEE/EMBS Conference on*, pages 1386–1389. IEEE, 2013.
- Xiao-Jing Wang. Neurophysiological and computational principles of cortical rhythms in cognition. *Physiological reviews*, 90(3):1195–1268, 2010.
- Zhijiang Wang, Zhengjia Dai, Gaolang Gong, Changsong Zhou, and Yong He. Understanding structural-functional relationships in the human brain: a large-scale network perspective. *The Neuroscientist*, 21(3):290–305, 2015.
- Takamitsu Watanabe, Satoshi Hirose, Hiroyuki Wada, Yoshio Imai, Toru Machida, Ichiro Shirouzu, Seiki Konishi, Yasushi Miyashita, and Naoki Masuda. A pairwise maximum entropy model accurately describes resting-state human brain networks. *Nature communications*, 4:1370, 2013.
- Takamitsu Watanabe, Naoki Masuda, Fukuda Megumi, Ryota Kanai, and Geraint Rees. Energy landscape and dynamics of brain activity during human bistable perception. *Nature communications*, 5, 2014.
- Aaron T Winder, Christina Echagarruga, Qingguang Zhang, and Patrick J Drew. Weak correlations between hemodynamic signals and ongoing neural activity during the resting state. *Nature neuroscience*, 20(12):1761, 2017.
- Fang-Chin Yeh, Aonan Tang, Jon P Hobbs, Pawel Hottowy, Wladyslaw Dabrowski, Alexander Sher, Alan Litke, and John M Beggs. Maximum entropy approaches to living neural networks. *Entropy*, 12(1):89–106, 2010.
- Benjamin E Yerys, John D Herrington, Theodore D Satterthwaite, Lisa Guy, Robert T Schultz, and Danielle S Bassett. Globally weaker and topologically different: resting-state connectivity in youth with autism. *Molecular autism*, 8(1):39, 2017.
- Mahinda Yogarajah, Niels K. Focke, Silvia B. Bonelli, Pamela Thompson, Christian Vollmar, Andrew W. McEvoy, Daniel C. Alexander, Mark R. Symms, Matthias J. Koepp, and John S. Duncan. The structural plasticity of white matter networks following anterior temporal lobe resection. *Brain*, 133(8):2348–2364, aug 2010. ISSN 1460-2156. doi: 10.1093/brain/awq175.
- Liang Zhan, Jiayu Zhou, Yalin Wang, Yan Jin, Neda Jahanshad, Gautam Prasad, Talia M Nir, Cassandra D Leonardo, Jieping Ye, Paul M Thompson, et al. Comparison of nine tractography algorithms for detecting abnormal structural brain networks in alzheimers disease. *Frontiers in aging neuroscience*, 7:48, 2015.
- Zhiqiang Zhang, Wei Liao, Huaifu Chen, Dante Mantini, Ju-Rong Ding, Qiang Xu, Zhengge Wang, Cuiping Yuan, Guanghui Chen, Qing Jiao, et al. Altered functional–structural coupling of large-scale brain networks in idiopathic generalized epilepsy. *Brain*, 134(10):2912–2928, 2011.