

Neural recordings and time series analyses

Randolph F. Helfrich and Robert T. Knight

Introduction

Modern day neuroimaging provides unprecedented spatiotemporal resolution to study cognitive processing at the level of whole-brain large-scale network interactions. For decades whole-head electroencephalography (EEG) was restricted to only 19 channels as defined by the international 10-20 electrode layout (a practice that prevails in the clinical setting). In contrast, modern day amplifiers enable simultaneous recordings of a few hundred channels (Biasiucci et al., 2019; Cohen, 2017; Lopes da Silva, 2013). For example, current high-density EEG systems offer up to 256 channels, which can be sampled at $>1,000$ data points per second. When combined with inverse solutions to project 2D sensor-level data into 3D source space, as informed by individual high-resolution structural MRI (magnetic resonance imaging) scans, one can obtain densely sampled time series data for $>4,000$ voxels inside the brain at 1 mm^3 resolution. A thorough discussion of inverse solutions exceeds the scope of this chapter, for a detailed discussion see (Lopes da Silva, 2013; Pascual-Marqui et al., 1994; Van Veen et al., 1997). This development now places high-density EEG recordings on par with magnetoencephalography (MEG; Baillet, 2017; Gross et al., 2013) in terms of spatiotemporal resolution and approaches the spatial resolution of functional magnetic resonance imaging (fMRI). This wealth of data poses a challenge for extraction of meaningful information about the functional architecture underlying human perception, cognition and action.

In this chapter, we focus on analysis techniques of time series data. First, we provide a brief overview of current methods that enable imaging the human brain with high spatial and temporal resolution. Throughout the chapter we emphasize that time series analyses can be applied to different types of electrophysiological data. Second, we review analyses strategies for high-dimensional time-series data. Methods are introduced according to their practical importance during data analysis (i.e., univariate analysis approaches in the time-domain are covered first, before advancing into spectral decomposition, bivariate connectivity analyses and finally multivariate analysis strategies). We then review methods that go beyond established linear time- and/or frequency analyses and discuss non-linear approaches, including information-theoretical approaches as well as recent machine-learning inspired strategies. Finally, we take recent developments of the last five years into account, as exemplified by strategies to analyze background ‘noise’, which has recently been shown to contain important behaviorally relevant information. In addition, we highlight how analysis strategies can be synergistically combined to maximize insight into neurophysiological processes underlying human cognition. Throughout the chapter, we highlight potential caveats, with the goal to provide a roadmap for state-of-the-art electrophysiological data analysis.

Neural recordings

Traditionally, neural recordings mainly referred to invasive single unit and local field potential recordings in animal models (Buzsáki et al., 2012). In the human literature, there is often a distinction between direct neurophysiological recordings, as for instance EEG, and indirect measures, such as the fMRI BOLD (blood oxygen level dependent) response. The goal of this chapter is not to provide an exhaustive list of imaging modalities, but rather to survey the analytical possibilities in the context of time-series analyses and highlight the similarities

between different methods. To accomplish this, we adopt a liberal definition of neural recordings, which encompasses every method that enables quantification of a brain process, irrespective whether it refers to direct electrophysiological recordings (voltage differences in scalp and intracranial EEG), magnetic fields (MEG), indirect oxygen-dependent responses (fMRI BOLD) or functional near-infrared spectroscopy; fNIRS) or motor signals that are obtained at the output stage, such as time-resolved behavior (quantified by time-resolved hit rates or reaction times), electromyography (EMG; voltage trace) or pupillometry and pupil size (recorded in millimeters, degree angles or pixels). These considerations can also be extrapolated to other types of recordings, such as two-photon calcium imaging data in rodents. One exception that we discuss in detail is that certain analysis approaches are limited by the sampling rate and recording location, such as extracting single unit responses from continuous local field potential recordings, which require recordings at a fine-grained spatial scale at high temporal resolution above 30,000 Hz.

To illustrate the analytical approaches, we focus on electrophysiological recordings in humans, both by means of scalp EEG as well as intracranial electrophysiology, which encompasses intracranial EEG, local field potentials and single unit activity (Parvizi and Kastner, 2018). While most readers will be familiar with scalp EEG, which will not be reviewed in detail (Biasiucci et al., 2019), we provide a more in-depth account of intracranial recordings in humans (Fried et al., 2014; Parvizi and Kastner, 2018).

Intracranial electrophysiological signals can be obtained from electrodes placed within the human brain for diagnostic and/or therapeutic purposes (Figure 1). The two most common entities that require a surgical placement of leads in the human brain are either placement to deliver therapeutic deep brain stimulation electrodes for Parkinson's disease (target: subthalamic

nucleus or internal globus pallidus; (Bronstein et al., 2011)), dystonia (ventral intermediate nucleus of the thalamus) or epilepsy (anterior nucleus of the thalamus; (Fisher et al., 2010)). In addition, implanted electrodes are used to guide surgical decision-making for intractable epilepsy (Parvizi and Kastner, 2018). Here, leads are inserted into multiple nodes of the suspected seizure network to identify the seizure onset zone. Target areas are identified according to the non-invasive work-up, which includes scalp EEG, high-resolution imaging and neuropsychology and can be complemented by various other diagnostic tools, including positron-emission tomography (PET), voxel-based morphometry (VBM) or MEG (Zijlmans et al., 2019). An important feature of some electrodes is that they feature a hollow lumen, enabling insertion of additional wire bundles, which protrude by 2-4 mm at the electrode tip for recording local field potentials and unit activity (Fried et al., 2014). Several studies have demonstrated the feasibility and safety of this approach (Carlson et al., 2018; Chari et al., 2020; Despouy et al., 2020; Hefft et al., 2013). Over the last decade, intracranial recordings in humans have yielded important insights into the functional architecture of cognition, such that the method is now widely regarded as an ideal tool to bridge the gap between invasive recordings in animal models and non-invasive recordings in humans.

Figure 1

Ultimately, only a few factors determine the application of time-series analysis strategies. Foremost are the sampling rate and the duration of the recordings. These factors determine the resulting Nyquist frequency (i.e., the highest frequency that can be resolved from the data is half the sampling rate – at 1,000 Hz sampling rate, all frequencies up to 500 Hz can be resolved; practically one should rather aim for a 3rd or 4th of the sampling rate). The duration further determines the Rayleigh frequency resolution, which is defined by 1 divided by the temporal

window. For example, a two second segment at 1,000 Hz can be analyzed in steps of 0.5 Hz (1/2), while a 10 second segment allows a more fine-grained resolution at 0.1 Hz (1/10). The signal duration also determines the digital filtering that can be applied to the data. The lower cut-off needs to match at the least one whole cycle of the lower boundary (i.e., filtering a 1 s second segment at 1 Hz is theoretically possible, but pushes the limits of signal processing leading to unstable results). However, filtering at a lower cut-off of 10 Hz (10 cycles in 1 second) is feasible, while a cut-off of 0.5 Hz is impossible (only half a cycle can be fit into the segment). Lastly, one needs to be aware of the noise floor of both the environment as well the recording equipment. For typical EEG amplifiers, the noise floor where the amplifier yields meaningful results is > 100 Hz. However, the environmental noise floor is already present at ~ 30 - 40 Hz given muscle activity in this recording range. This issue is largely mitigated in intracranial EEG recordings, which are less impacted by muscle artifacts. For instance, analyses in the high-frequency activity band (70-200 Hz) have yielded important insights into cognitive functioning (Leszczyński et al., 2020).

Keeping those theoretical principles in mind now enables applying spectral analyses to different recording modalities with theoretical and practical implications (Prerau et al., 2017). For example, intracranial EEG data that was recorded over one hour at a frequency resolution of 5,000 Hz can easily be filtered and spectrally decomposed in a broad-range of frequencies up to 2,500 Hz at a fine-grained resolution of $\ll 0.1$ Hz (in practice an upper threshold of 250 Hz and a resolution of 0.5 Hz is often sufficient). In contrast, whole-brain fMRI BOLD over one hour typically provides one data point per voxel every second, thus, resulting in a frequency resolution of 0.33 Hz, which limits both the ability to filter as well as to spectrally decompose the data (Fox et al., 2005). For fMRI using these parameters, the upper frequency cut-off is at 0.16 Hz, so

frequencies between 0.01 and 0.1 Hz can best analyzed at a resolution in the 10^{-2} to 10^{-3} Hz range.

Electrophysiological data analysis

The improved spatiotemporal resolution of human electrophysiological recordings has benefited the development of numerous algorithms and methods to extract behaviorally meaningful information. However, these high-dimensional and complex data sets that are governed by non-linear dynamics, posing unique challenges. For example, there are multiple correct answers for the questions regarding (1) how to analyze the data, (2) which method to choose and (3) how to statistically quantify the results. More commonly, the right answer will be ‘it depends’. Here, we provide a practical road map with an initial focus on EEG analyses that helps narrow the immense search space and justify analytical choices and their interpretations.

General analysis strategies

Data analysis does not begin once data collection ended, but rather starts with data recording and the experimental structure. During data recording, researchers make a number of explicit (e.g., number of EEG sensors, sampling rate) as well as implicit choices (task design, duration of trials, event structure) that limit the subsequent analyses. Important considerations include whether distinct events are present (task-based time-locked events that can be contrasted) or whether a continuous design (e.g., resting state or experience sampling) were employed. Within the historical context of EEG, Hans Berger’s seminal experiments on alpha oscillations (cf. Berger, 2004) first constituted a continuous analysis (spontaneous fluctuations in the trace were quantified). Subsequently, he employed an event-locked approach (eyes open/close), which later gave rise to the discovery of event-related potentials by averaging across several repetitions (Polich, 2007). To illustrate these methods, we will assume a task-based design with distinct

events that require a behavioral response and thereby, enable time-locked analyses either relative to the event or the action.

Event-related analysis in the time-domain

Averaging across several events is typically done to improve the signal-to-noise ratio and to quantify the average response across a number of events. In the context of EEG, averaging can be done on the raw trace, thus, yielding event-related potentials (ERP; Handy, 2005)). Likewise, averaging can be performed after extraction of a distinct spectral component. For example, it had been observed that the 70-150 Hz (termed high gamma or high frequency band activity, HFA/HFB) in intracranial EEG recordings reflects a proxy of multi-unit activity and contains more behaviorally relevant information than the broadband signal (Edwards et al., 2005; Flinker et al., 2015; Kanth and Ray, 2020; Leszczyński et al., 2020; Ray and Maunsell, 2011; Rich and Wallis, 2017). Thus, filtering and extracting the signal envelope by means of (e.g., a Hilbert transform; Figure 2), are used to distil the relevant signal component prior to subsequent averaging (Figure 2A).

Figure 2

In continuous task designs, such as resting state or sleep recordings, there are no external temporal structure for time locking analyses. Therefore, researchers rely on identification of intrinsically generated events for subsequent analysis. Many events have first been described in the time domain (e.g., alpha waves by Berger as characteristic 10 Hz bursts of activity in the raw trace). Likewise, slow waves (< 4 Hz) and sleep spindles (~12-16 Hz) have been identified based on the characteristic waveform shape in continuous recordings (Buzsáki, 1996; Diekelmann and Born, 2010). Detection of these events in the time domain can be algorithmically formalized, typically by introducing constraints with respect to (a) the frequency content (through band pass

filtering), (b) an amplitude criterion, (c) a duration criterion and (d) exclusion of unrelated activity (Helfrich et al., 2019; Staresina et al., 2015). For example, detection of sleep spindles only considers bursts in the amplitude series of the 12-16 Hz filtered signal that exceed a z-score of (e.g., 3 SDs above the signal mean) for anywhere between 0.5 to 3 seconds. Based on the individual detections, averaging can be performed in the time domain to obtain the average waveform shape (Figure 3).

Figure 3

Furthermore, the detected event can be conceptualized as a point process, i.e. an event that occurred precisely at one moment in time (e.g. activity peak). To again employ the example of sleep spindles, time-locking of the raw signal relative to spindle peaks reveals the presence of a second spectral signatures in the raw signal. The spindle does not occur in isolation, but is nested with a slow wave. This approach exemplifies how detection of endogenous temporal events can reveal temporal regularities that otherwise could not be detected from the raw trace.

The same principle applies to the extraction of single unit spikes from broadband local field potentials (LFP). LFPs are sampled at $>30,000$ Hz, spikes are again extracted based on amplitude and time criteria (sharp transient in the range from 2-4 ms) before all waveform shapes are further characterized in a process called spike sorting (Fried et al., 2014; Rutishauser et al., 2006). Here, different waveform shapes are disentangled in order to isolate activity from distinct neurons. Subsequently, every spike from every identified neuron is regarded as one distinct time stamp; thus, the dimensionality of the data has been reduced, benefiting subsequent analyses. For example, spike-triggered averaging of the raw broadband signal has revealed that spikes preferentially occur at distinct phases of the underlying population oscillation, which are already present in the down-sampled signals. Hence, subsequent analyses can be carried out on

signals sampled at 1000 Hz and not 30,000 Hz, enhancing computational efficacy and speed. Furthermore, from these point events ('Spikes'), peri-stimulus time histograms (PTSH) can be created, which are commonly smoothed into a trace of overall spiking activity; hence, yielding a signal that is comparable to an event-related potential, only based on a number of distinct neuronal spikes. Collectively, this demonstrates how temporal events can be extracted from continuous data and then again be analyzed in a time-locked fashion, either relative to external or internal events.

Spectral analysis

Electrophysiological signals are rich and complex. Ever since Berger's seminal observation, the community was well aware that different frequency bands might contain distinct information. Over recent decades, several methods were introduced to either spectrally decompose the signal to obtain activity in multiple frequency bands or to isolate activity in a distinct frequency band.

The most common approach is based on the Fourier transformation, which decomposes the signal into (co-)sines and provides estimates of the relative contribution of every frequency band to the entire signal. Using this approach, the time domain is typically lost (translation of time- into frequency-domain), but can be recovered by means of repeating the analysis in different time windows by means of convolution. Electrophysiological power spectra exhibit distinct characteristics, including a steep $1/f^x$ drop-off, where x typically scales in the range from -2 to -4 in the healthy brain (for extended discussion see section on aperiodic activity; (He et al., 2010; Miller et al., 2009)). Band-limited oscillations arise as distinct 'bumps' (Donoghue et al., 2020) above the background activity (Figure 4).

Figure 4

In addition to the Fourier transform, several other methods have been introduced, such as Morlet wavelets (sines multiplied with Gaussians, which are convoluted with the signal) or the Hilbert transform, which requires band-pass filtering first and then estimates the instantaneous amplitude and phase series. One important caveat is that all spectral decomposition techniques only provide ‘estimates’, which differ as a function of (e.g., signal length, sampling rate, noise and pre-processing). A common approach to improve estimates is to ‘window’ the data (also called taper; multiplication of a data segment with (e.g., Boxcar window, where the edges are Gaussian shaped, thus, attenuating edge artifacts). The signal-to-noise ratio can be improved if multiple windows are combined and estimates are averaged (as for instance done when using Welch’s method or the multi-taper approach based on discrete prolate spheroidal sequences (Prerau et al., 2017); Figure 5).

Figure 5

Another major drawback that all methods share is that they are based on sinusoidal basis functions (i.e., if these methods are applied to brain data, then the results will be systematically biased to reveal sinusoidal oscillations). Inspection of raw EEG traces reveals that most neuronal oscillations are not sinusoidal in nature, but are characterized by biased rise- and decay-times with skewed and often sharp waveform shapes (Cole and Voytek, 2017). Application of sinusoidal methods can introduce severe artifacts, which are prone to misinterpretation (Aru et al., 2015; Gerber et al., 2016). One example is sensorimotor Mu-Rhythm at ~8-12 Hz, which is named ‘Mu’ given that it is often shaped like a ‘M’. The Fourier transform will extract the main component at ~10 Hz, but also yield peaks at ~20 Hz as well as all other subsequent harmonics, given that the sharp peak is incompletely captured by a single sine wave (Voytek et al., 2010). In electrophysiological recordings, researchers are then faced with the presence of true beta-band

activity at ~20 Hz (a hallmark of the motor system), which is contaminated by Mu-harmonics at ~20 Hz (Stolk et al., 2019). In recent years, several methods have been introduced to disentangle true from spurious oscillatory brain activity (Donoghue et al., 2020; Kosciessa et al., 2020; Wen and Liu, 2016). For example, empirical mode decomposition (EMD) does not rely on sinusoidal basis functions and can also model a change in the precise peak frequency over time (Quinn et al., 2021). The same issues apply when the interaction of multiple frequency bands is assessed. To date, there is no unique solution for this issue, but several recent publications suggest that a set of criteria could be applied to the data to infer whether non-sinusoidality is present (Aru et al., 2015).

Separating aperiodic from oscillatory activity

The electrophysiological power spectrum encompasses oscillations (discrete ‘bumps’), exceeding the general $1/f^\alpha$ drop-off. For decades, both phenomena have been studied together and were not explicitly disentangled (Donoghue et al., 2021). Hence, elevated ‘alpha power’ could either be the results of an amplitude increase of the oscillatory component or a general increase in activity in all frequency bands (Figure 4). This distinction has gained more traction recently, since several groups demonstrated that the background activity – which had been previously considered to mainly reflect neuronal noise – in fact contains behaviorally relevant information (Donoghue et al., 2020; He et al., 2010; Lendner et al., 2020; Voytek et al., 2015). Importantly, the level of information content about the behavioral state is on par with neuronal information encoded in band-limited oscillatory activity. To date, the exact physiological role of both components is not fully understood, but they can be conceptualized as providing complementary insights into cognitive processes (Wainio-Theberge et al., 2021, 2022). Over the last decade, several approaches have been introduced to parameterize oscillations and to

disentangle oscillatory from broadband (also termed aperiodic, fractal or scale-free) background activity. The simplest solution to isolate oscillations entails a linear fit to the power spectrum in double-log space (thus the linear fit reflects the aperiodic component) and subtraction from the original spectrum (Lendner et al., 2020; Miller et al., 2009). More sophisticated algorithms employ robust fitting (Kosciessa et al., 2020), irregular spectral resampling (Wen and Liu, 2016) or additive fitting of Gaussian components to the oscillatory peaks (Donoghue et al., 2020). All methods have in common that they now provide two components that are derived from the same underlying signal, which can be related to behavior and brain state. Currently, these concepts are actively being explored, with several new notions, such as the concept of population time constants of relative stability (also termed temporal integration windows, autocorrelation window or intrinsic neural timescales; Gao et al., 2020; Golesorkhi et al., 2021; Ito et al., 2020; Raut et al., 2020; Wolff et al., 2022), which can be approximated by characteristic bends in the shape of the power spectrum in the frequency domain or from the decay of the autocorrelation function in the time domain. Population time constants are thought to provide the necessary means for temporal integration and exhibit a clear cortical gradient with short timescales in sensory cortex and longer timescales in association areas (Gao et al., 2020; Raut et al., 2020; Wolff et al., 2022).

Oscillatory and waveform shape features

Definition of oscillations, including their bandwidth and amplitude (i.e., power), is most commonly done in the spectral domain. However, as outlined above, spectral analysis also omit important wave form features as well as instantaneous signal characteristics, which are inherently time dependent, such as oscillatory phase. Hence, it is best practice to return from the spectral to the time domain to quantify and assess the oscillatory features after the presence of an

oscillation was established using spectral methods. These analyses can either be applied on band-limited or broadband data. It is important to consider that both linear as well as non-linear variables can be obtained from time- and frequency-domain data, with amplitude reflecting a linear variable, while phase constitutes a circular, and thereby, non-linear variable. Both features can be extracted from the Hilbert transform and can be related to behavior using either linear or circular-linear correlations (Berens, 2009; Fiebelkorn et al., 2018; Helfrich, Fiebelkorn, et al., 2018). In this case, one needs to assume a fixed frequency band to obtain reliable phase estimates. Again, Fourier methods are ill suited to assess if the peak frequency changes over time (jitter in peak frequency reflects a broadening of the spectral peak in the electrophysiological power spectrum). Several recent developments, including the EMD, now aim at mitigating these effects (Quinn et al., 2021; Watrous and Buchanan, 2020).

Another recent development is the appreciation of waveform shapes (Cole et al., 2017). Previously, waveform shapes have largely been ignored given that the Fourier basis functions were sinusoidal. The role of physiologic non-sinusoidal waveform shapes came into focus after it was realized that non-sinusoidality introduces artifacts in (e.g., cross-frequency coupling analyses; see below; Aru et al., 2015; Gerber et al., 2016); Figure 6). Novel tools now enable assessing waveform shapes with the goal to relate distinct features (such as rise- or decay-times, asymmetries or amplitude bias) to distinct physiological processes. However, there is currently only limited evidence that supports a distinct role in cortical processing.

Figure 6

Network connectivity across spatial and temporal scales

We have focused on univariate analyses that can be carried out at the single electrode or single voxel level. In contrast, network approaches quantify the interaction between multiple regions.

This interaction is often also termed functional connectivity, to contrast it to structural connectivity, which can be obtained from tracing or fiber tracking studies (Buzsáki and Draguhn, 2004; Engel et al., 2001; Varela et al., 2001). Functional connectivity is undirected (i.e., it is not quantified who drives the interaction). Directed connectivity is referred to effective connectivity implying information flow from node A to node B. Here, we first discuss undirected connectivity between different nodes, and then we discuss directed connectivity methods across spatial scales. The concept of connectivity can also be extended to interactions in the temporal domain (i.e., the interaction between different oscillations or between the LFP and spikes. Lastly, we discuss the use of information theoretical approaches in this context.

Undirected connectivity is most commonly studied by means of coherence or magnitude-squared coherence as it is more correctly termed (Bastos and Schoffelen, 2015). This distinction indicates that the coherence formula encompasses both the relationship of amplitude as well the relationship of phase. Collectively, this relationship is then normalized to yield a number between 0 and 1. If the amplitude term in the formula is replaced by a 1, then the formula become amplitude-independent and only the contribution of phase synchronization is estimated; this normalized variant is also known as the phase-locking value (Lachaux et al., 1999). Likewise, one can focus only on the amplitude contribution by (e.g., correlation of the amplitude time series; Hipp et al., 2012). The literature often distinguished between phase-based and amplitude-based connectivity and several theoretical accounts postulated distinct roles for cortical communication (Engel et al., 2013). Both metrics have in common that they are susceptible to volume spread in the cortical tissue, which inflates connectivity metrics. Solutions by means of orthogonalized amplitude correlations (Hipp et al., 2012) or imaginary coherence (Nolte et al., 2004) as well as several variants have been proposed (Bastos and Schoffelen,

2015), which minimize this confound by attenuating the contribution of zero phase-lag interactions.

Both approaches share that they are undirected in nature, so it remains equivocal whether node A is driving node B, or vice versa. Several metrics that take advantage of statistical regularities in the data have been introduced to infer directionality, such as Granger causality (model-based assessment; (Seth et al., 2015)) or the phase slope index (dependence across multiple frequency bands; (Nolte et al., 2008)). However, these methods operate on statistical dependencies, are susceptible to noise and do not provide a true ‘causal’ explanation.

The concept of directed and undirected connectivity has also been formalized in an information-theoretical framework (Ince et al., 2017; Panzeri et al., 2015). Shannon Information Theory (Shannon and Weaver, 1998) is centered on entropy to quantify the observed distribution of a given variable (here activity in one region). Connectivity between two regions can be inferred if knowledge about the activity in one region reduces uncertainty about the state of the other region. This interaction has also been termed Mutual Information and is undirected. An important extension of this idea called sample entropy basically takes a third signal into account to infer directionality (Lobier et al., 2014): Does knowledge about the past of region A provide more information about the future of region B than the past of region B alone? If the answer is yes, then this can be interpreted as information flow from A to B. Again, this approach is not causal, but enables a statistical comparison based on empirical signal distributions. One disadvantage of information theoretical metrics is that they typically require binning of the data and thus, reduce the signal complexity into a finite number of bins at the expense of losing fine-grained details (Panzeri et al., 2015).

Lastly, the concepts that apply to interactions between spatial nodes can be extended to the spectral domain to infer if two frequency bands interact (Figure 7). This coupling across temporal scales is widely known as cross-frequency coupling (CFC; (Canolty and Knight, 2010) and has most prominently studied in the context of phase-amplitude coupling (PAC), where the phase of slow frequency (e.g., theta or alpha) predicts broadband high frequency activity (Canolty et al., 2006). However, amplitude-amplitude as well as phase-phase couplings have also been described for CFC (Aru et al., 2015). CFC analyses are susceptible to signal artifacts (cf. Figure 6) and signal processing choices; hence, several papers formulated concrete guidelines to circumvent these issues (Aru et al., 2015; Gerber et al., 2016). Similar considerations apply also for LFP-spike coupling.

Figure 7

Multivariate representations and analyses

The analysis of electrophysiological data constitutes an extensive multivariate problem. In order to better understand the data and help the interpretation, the most common analytical approaches condense this complex analytical space into serial univariate tests, which are oftentimes easier to interpret. From a conceptual standpoint, univariate analyses seem reasonable when data is approached with the ‘Neuron Doctrine’ in mind, which states that the single neuron is the central computational unit of the nervous system (cf. Yuste, 2015). Within this framework, it is reasonable to employ univariate analyses approaches. However, in recent years it became obvious that the single neuron framework falls short in explaining several cognitive phenomena, such as flexible context-dependent behavior, which inspired the concept of a ‘Population Doctrine’ (Ebitz and Hayden, 2021; Eichenbaum, 2017; Saxena and Cunningham, 2019; Yuste, 2015). This population framework suggests that transient coalitions of neurons form the central

computational unit of the nervous system (Eichenbaum, 2017; Siegel et al., 2012). Therefore, gaining a deeper understanding into cognitive mechanisms requires large-scale recordings, which are not amendable to classic univariate analyses. This novel framework was motivated by a series of new artificial-intelligence inspired algorithms, which enable uncovering the organizing principles of seemingly chaotic population activity. This is a rapidly emerging field with many new technical innovations in recent years, and we are unable to provide a thorough review, but will focus on two representative examples to illustrate how to approach population-based analyses.

The first important approach is that multivariate data can be represented in a so-called state-space representation (Ebitz and Hayden, 2021; Gervasoni et al., 2004). Instead of analyzing individual time series, the entire channel (or neuron) x time matrix is conceptualized as a single time series of adjacent points which travels through a N-dimensional coordinate system along a trajectory, which is spanned by the individual data points. Here N refers to the number of observations (e.g., channels or neurons). If one considers recordings from three neurons, then the first time point is a point in a 3D coordinate system. Hence, using linear algebra, we can infer the Euclidean distance to either the center point of coordinate system or to adjacent time points. Here the critical advantage is that Euclidean distance is again a single number, hence, the high dimensional data can effectively be condensed into a single vector. Thus, the entire matrix can be regarded as a trajectory traveling through N-dimensional space and given that Euclidean distances can be calculated, one can also infer (e.g., the velocity (distance/time) or acceleration (velocity/time) of the system — Figure 8).

Figure 8

As humans, we cannot effectively visualize more than three dimensions, hence, it is common practice to visualize state space trajectories in 2D or 3D, and after dimensionality reduction has been performed (Ebitz and Hayden, 2021). This conceptualization of population activity gave rise to several emerging concepts in neuroscience, including attractor states (points in the system where the activity patterns are naturally drawn to and converge on) or neuronal manifolds (activity patterns fall onto certain planes in the state space and are not completely randomly distributed along all possible dimensions, hence, all possible activity patterns are constrained to a subspace that houses most observations). To date, these techniques have mainly been applied in the motor domain, where they provide additional explanatory power over univariate analyses, but they are gaining traction in the cognitive domain (Chaisangmongkon et al., 2017; Goudar et al., 2021; Murray et al., 2017; Weber et al., 2021). The second common technique to address multivariate brain data relies on multivariate pattern analysis (MVPA; (Grootswagers et al., 2017; Hebart and Baker, 2018; Kriegeskorte et al., 2006; Quian Quiroga and Panzeri, 2009), which is also sometime called pattern classification or more commonly ‘brain decoding’ (Figure 9).

Figure 9

Here, researchers take advantage of classification algorithms, such as the linear discriminant analysis (LDA) or support vector machines (SVM), which learn patterns that are associated with certain conditions or responses. Importantly, the classifiers require splitting the data into a training dataset, where the classifier has access to the ground truth, and a testing dataset, which consists of held-out samples. Based on the performance on the held-out samples, one can infer the accuracy of the classifier. Critically, these algorithms often work as a ‘black box’, where the algorithm learns a given association and then is used to predict held-out data

points, but one cannot visualize the classifier per se. Therefore, it often remains ambiguous whether the classifier utilizes information that is also accessible by the brain itself or whether it picks up epiphenomenal or noise influences. In the context of neurophysiological experiments, classifiers are used to infer if a population contains behaviorally relevant information or not, i.e. whether the classifier performs better than chance. In contrast to engineering, the goal is not to maximize classifier performance, but the goal is to test if the classifier performs significantly better than chance. If above-chance performance is observed, this is taken as evidence that the brain contains information about the classified behavior (i.e., context or movement execution).

In practice, state-space and classification analyses are often combined (Mante et al., 2013). Jointly these methods enable extraction of coding dimensions (i.e., finding a latent or hidden dimension in the data that cannot be accessed using univariate analyses, that maximizes a given behavioral contrast). This approach enables disentangling (i.e., cognitive from motor contributions, which can appear superimposed when using univariate analysis tools; Vyas et al., 2020). Disentangling the respective contributions provides the opportunity to study their contribution to the overall population activity and provides new insights into the coding mechanisms. For example, it had repeatedly been shown that motor activity is confined to a low-dimensional sub-space (i.e., can be described by only few activity patterns, which often cycle along a low-dimensional ring structure; Churchland et al., 2012; Shenoy et al., 2013). In contrast, more complex higher-order cognitive operations also require higher dimensional neuronal representations, since they rely on distributed computing (Ebitz and Hayden, 2021). Collectively, novel population-based analysis strategies extend the explanatory power of univariate tests and open a new conceptual space to interpret the functional architecture underlying human behavior.

Scale-free phenomena

Scale-free phenomena refer to the analysis of statistical regularities in electrophysiological data that are not governed by a defining temporal scale (He et al., 2010). For example, scalp EEG is governed by alpha oscillations at ~ 10 Hz as their defining temporal scale, while many intracranial EEG contacts often do not exhibit a clear oscillatory peak. The absence of oscillatory peaks raises the question if there is any underlying hidden temporal structure present and if yes, then how could this be quantified. An important concept that borrowed from physics is fractality, which describes the self-similarity of a signal irrespective of the chosen temporal window (Pritchard, 1992). Critically, this self-similarity follows a power-law (i.e., a law of nature; Miller et al., 2009). A typical example from nature is the structure of a Romanesco cauliflower, which exhibits characteristic florets, which always look similar, irrespective of their size or magnification (Hardstone et al., 2012; Linkenkaer-Hansen et al., 2001). In the context of electrophysiological data, this means that an EEG trace always exhibits similar fluctuations, irrespective of whether one hour, one minute or one second is displayed. While this phenomenon is difficult to describe in the temporal domain, it can easily be captured in the frequency domain (Figure 10).

Figure 10

After applying a Fourier transform, one can easily appreciate the general $1/f$ drop-off (cf. aperiodic activity in EEG), which indicates a stable relationship between frequency and amplitude, irrespective of the precise frequency and therefore irrespective of the precise temporal scale. In the past two decades, several methods have been introduced to quantify this type of behavior, either in the temporal or spatial domain. For example, linear fitting of the $1/f$ background activity is a common approach to extract the spectral exponent and thereby, quantify the underlying statistical regularities. A related approach is the detrended fluctuation analysis

(Hardstone et al., 2012), which assesses long-range temporal correlations, i.e. signal correlations that obey the same scaling behavior irrespective of the chosen temporal window. Similar 1/f phenomena have been observed in the spatial domain, where activity peaks across many electrodes coincide in time and the distribution of the duration and amplitude again follow a power law with 1/f-like characteristics (Palva et al., 2013). Going beyond the activity peaks, analysis of microstates (i.e., quasi-static scalp topographies in EEG) revealed that brain activity alternates between four major scalp topographies, where again a 1/f-like power law can describe the duration of each topography (Ville et al., 2010). To date, it remains unclear how these different phenomena are related, but several lines of research suggest that they might constitute emergent analytical tools to quantify temporal regularity, which cannot be appreciated by the naked eye, unlike neuronal oscillations (Cocchi et al., 2017). It has been argued that the characterization of temporal regularities in the form of scale-free dynamics constitutes a promising avenue to understanding the neural correlates of consciousness (Lendner et al., 2020; Tagliazucchi et al., 2013, 2016; Zhang et al., 2018).

Statistical considerations

Since brain activity is non-linear, non-stationary and non-normally distributed, classic analyses such as t-tests, analysis-of-variance (ANOVA), linear correlation or regression often fall short in capturing the data complexity (Maris and Oostenveld, 2007). However, these methods are still widely employed, since they are easy to use and most researchers know how to interpret their outcomes (Piai et al., 2015).

To mitigate the issue of unknown distributions of brain data, which violate the assumptions of many statistical tests, the currently most accepted approach includes cluster-based permutation tests (Maris and Oostenveld, 2007). Here a new surrogate distribution is build

based on randomly assigning all observations to distinct groups (i.e., shuffling the condition label), then to repeat the analysis 1,000-10,000 times and then compare the true observed value to the surrogate distribution. If the observed value falls outside of the mean ± 2 SDs, then significance can be assumed (at $p < 0.05$). This approach is elegant, since it allows that any type of primary test statistic can be employed to build the surrogate distribution. Hence, this test is inherently non-parametric and can be used in a wide variety of contexts. An alternative solution that emerged in recent years, is the use of general linear mixed effects models, which however, are often difficult to interpret and difficult to visualize and therefore, have not been adopted widely.

Summary and Conclusions

In the present chapter, we have reviewed the current state-of-the-art of neuronal time series analysis through the lens of scalp and intracranial EEG recordings. Notably, as outlined in beginning, the same analyses can be applied to other imaging modalities that produce time series data and most analyses are only constrained by sampling rate, signal duration and event types. From a practical standpoint, we emphasized the need to go from simple to more complicated analyses, i.e. going from univariate analyses in the time domain to univariate analyses in the frequency domain, before branching off into bivariate connectivity or multivariate classification analyses. In order to interpret the result of a bivariate connectivity or multivariate classification analysis in a meaningful way, one needs to ensure that the observed differences are not solely driven by univariate differences, in (e.g., the overall amplitude), but in fact can only be attributed to their joint activity.

The example of aperiodic activity analyses nicely illustrates how one can teach an old dog (the Fourier transform) a few new tricks that provide additional explanatory power and

enable linking macroscopic signals to underlying circuit properties (Gao et al., 2017). In the case of aperiodic activity, computationally modeling suggested that the steepness of the exponential decay of the power spectrum reflects the excitation-inhibition-balance of the underlying neuronal population, which otherwise could not have been inferred (Chini et al., 2021).

The wealth of available methods also enables reanalysis of available datasets, especially if the primary analysis was mainly concerned with univariate analysis approaches. As an example, a classic working memory experiment in two monkeys indicated that single neurons in prefrontal cortex code distinct aspect of working memory (Warden and Miller, 2007, 2010). Through careful reanalysis of the same dataset, a new concept of working memory emerged. In the first follow-up publication, the authors demonstrated that spiking is coupled to the underlying LFPs (Siegel et al., 2009). Subsequently, these LFP signatures were described in more detail, which triggered follow-up investigations on neuronal bursts and how these differ from sustained oscillatory activity (Lundqvist et al., 2016, 2018). Furthermore, the same dataset was used to infer the dimensionality of spiking activity, which supported the notion that high-dimensional representations are desirable for cognitive flexibility (Rigotti et al., 2013). With the emergence of more analytical tools and the availability of more open datasets novel concepts will be generated based on existing data (Fusi et al., 2016). Overall, machine learning is now established as an independent field, thus, we foresee that more multivariate tools will become available to experimental neurophysiologists to assess their data beyond classical univariate analyses (Grootswagers et al., 2017; Hebart and Baker, 2018).

Overall, experimentalists are nowadays faced with a wide-array of analysis tools, which requires careful justification of the employed methods. However, the wealth of tools also

provides an exciting avenue to unravel the mechanisms behind human cognition, perception and action.

Acknowledgements

This work was supported by the German Research Foundation (Emmy Noether Program, DFG HE8329/2-1; R.F.H.), the Hertie Foundation (Network for Excellence in Clinical Neuroscience; R.F.H.), the Baden Württemberg Foundation (PostDoc Program; R.F.H.), the Ernst Jung Foundation for Science and Research (Career Development Award; R.F.H.) and NINDS grants R01 NS021135 and U19NS107609-01 (R.T.K.).

Figures Captions and Legends

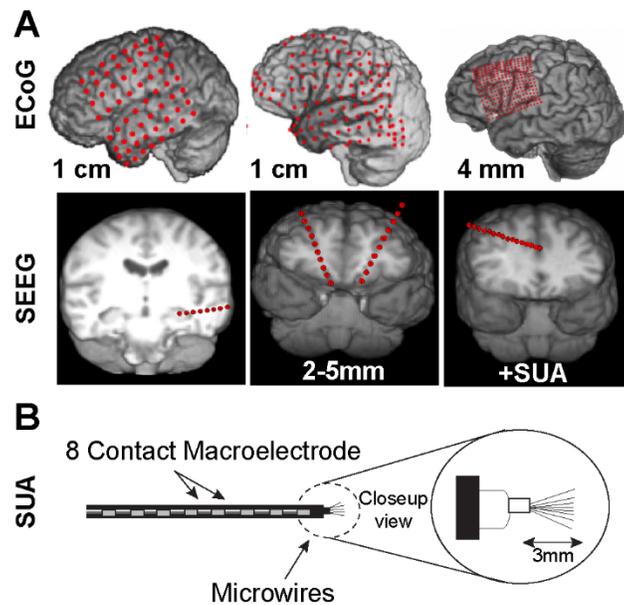


Figure 1. Intracranial human electrophysiology

(A) Upper row: Subdural grid electrode placement from three representative subjects covering large portions frontal, temporal and motor areas. Lower row: Stereotactically placed depth electrodes targeting medial temporal, orbitofrontal or medial frontal structures. (B) Single unit activity can be recorded from additional wire bundles that are inserted through a hollow lumen of the clinical macro-electrode. Figure reproduced with permission (Helfrich & Knight, 2019).

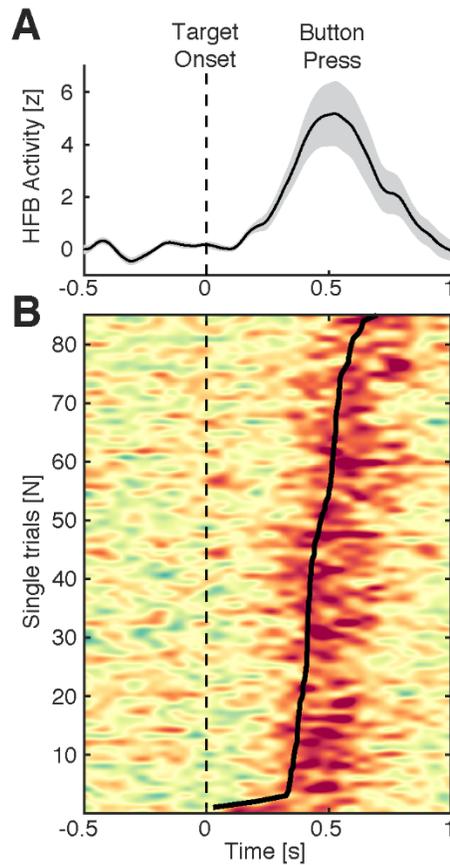


Figure 2. High frequency band activity

(A) Average high frequency band activity recorded from a motor cortex electrode relative to the detection of the onset of a visual target reveals activity after approximately 500ms. (B) Stacked single trials that are sorted relative to reaction times (black trace) reveals that HFA tracks behavior on a single trial basis. Unpublished (Helfrich & Knight).

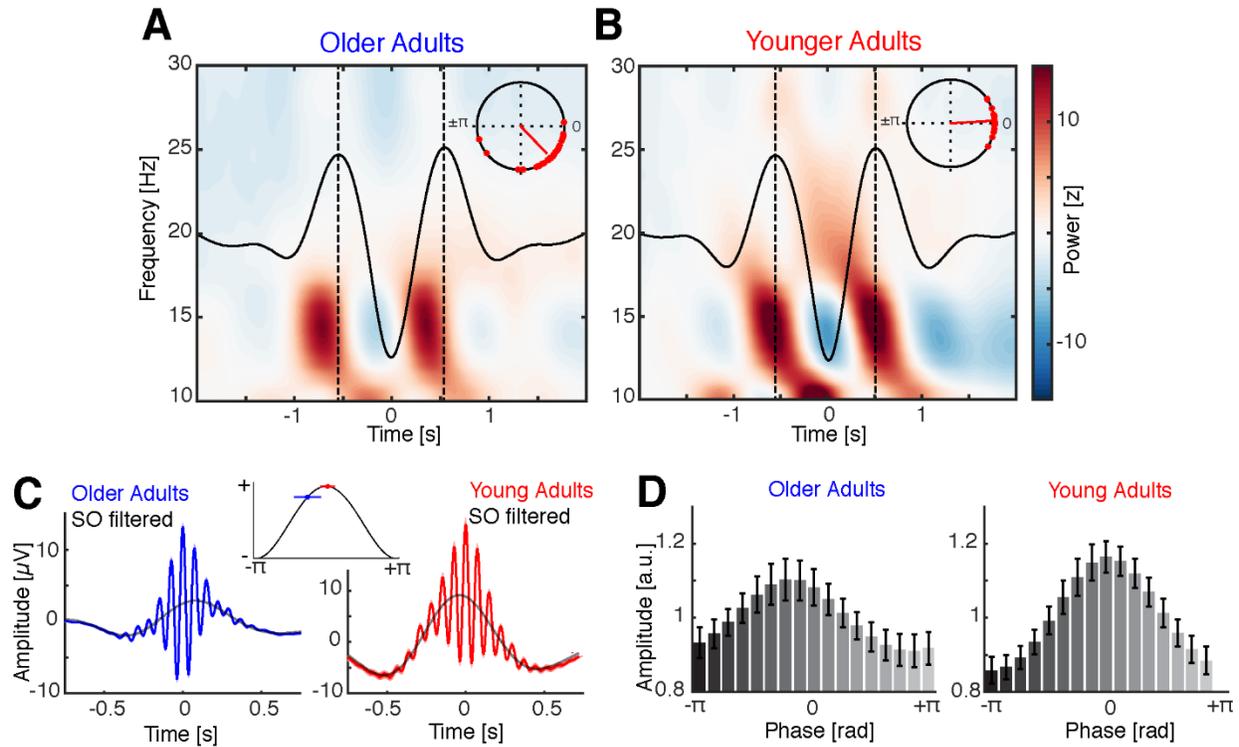


Figure 3. Visualization of cross-frequency coupling and raw waveform shapes

(A) Slow waves in a group of older adults (black) on top of the respective time-frequency decomposition, indicating an activity peak in the ~14 Hz spindle range just prior to the slow wave peak. Inset: Visualization in circular space of the same data, indicating that spindle activity peaked prior to the slow wave peak (approx. 45° , corresponding to 50-100ms). (B) Same visualization for a cohort of younger adults, indicating that spindles are precisely locked to the peak of the slow wave. (C) Visualization relative to detected spindles (colored) with the slow wave superimposed in black. Note that spindle peak prematurely. Inset: Illustration of the spindle peak relative to the slow wave (as exemplified by a cosine). (D) Histogram-based visualization of coupling: Relative amplitude in the spindle-band relative to different binned slow-wave phases, reveals a non-uniform distribution that is skewed in older adults. Note all panels depict the same data and illustrate different approaches to cross-frequency coupling and waveform shape analyses. Figure reproduced with permission from (Helfrich, Mander, et al., 2018)

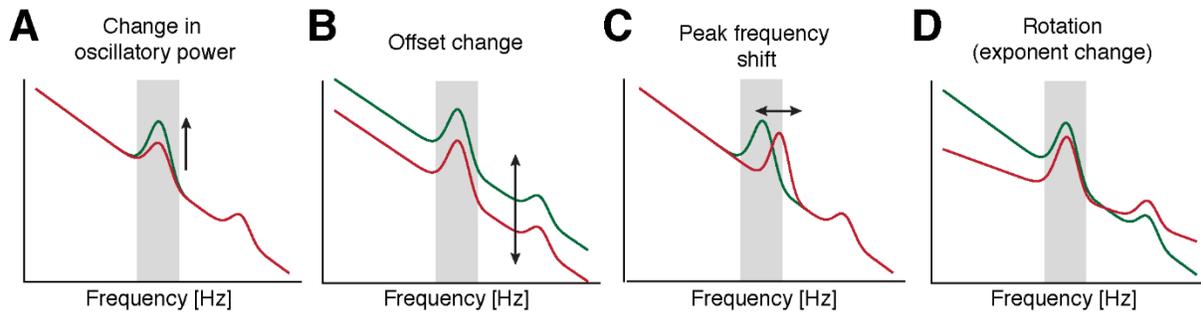


Figure 4. Relationship of oscillatory and broadband activity

(A) Illustration of Fourier transformation of electrophysiological data that exhibits both an oscillation in 8-12 Hz alpha range (‘bump’ exceeding the 1/f background activity within the gray shaded area) as well as broadband 1/f component. A true increase in oscillatory power between two conditions (red to green) is illustrated. (B) A broadband shift (i.e., change of overall offset along y-axis) can mimic an alpha power increase, which is not oscillatory in nature. (C) Likewise, a change in peak frequency can mimic power changes when power is averaged within a predefined range. (D) Similarly, a rotation of power spectrum (change in spectral exponent) can introduce changes in oscillatory power if spectra are not properly parameterized. Illustrations created using the FOOOF toolbox (Donoghue et al., 2020).

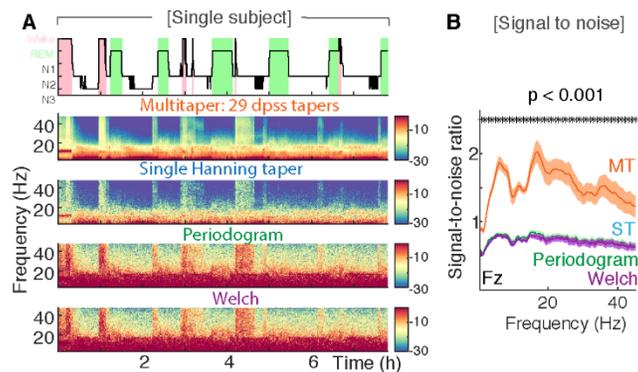


Figure 5. Effect of different spectral decomposition methods

Upper row: Sleep hypnogram from single case across 8 hours of sleep. Below: Spectral decomposition of entire night of sleep using either a multi-taper or single taper approach, the periodogram or Welch’s method. The resulting estimates become progressively more coarse-grained. (B) Signal-to-noise ratio per frequency reveals a clear advantage of multi-tapered spectral analysis. (from Lendner et al., 2020).

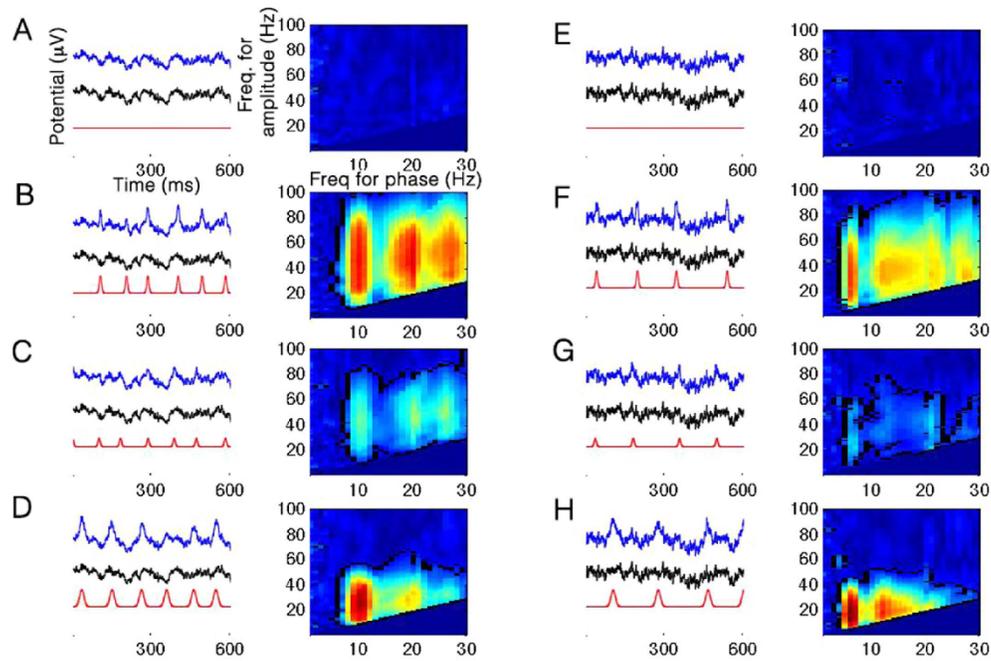


Figure 6. Effects of non-sinusoidal sharp transients on cross-frequency coupling estimates (A-H) Eight examples of simulated EEG traces (blue) that are composed of broadband EEG activity (black) and sharp transients (red) as well as the accompanying phase-amplitude comodulogram (black outlines indicate significant coupling) revealing that sharp transients introduce spurious cross-frequency coupling at the primary frequency (10 Hz) as well as at the subsequent harmonics. (from Gerber et al., 2016).

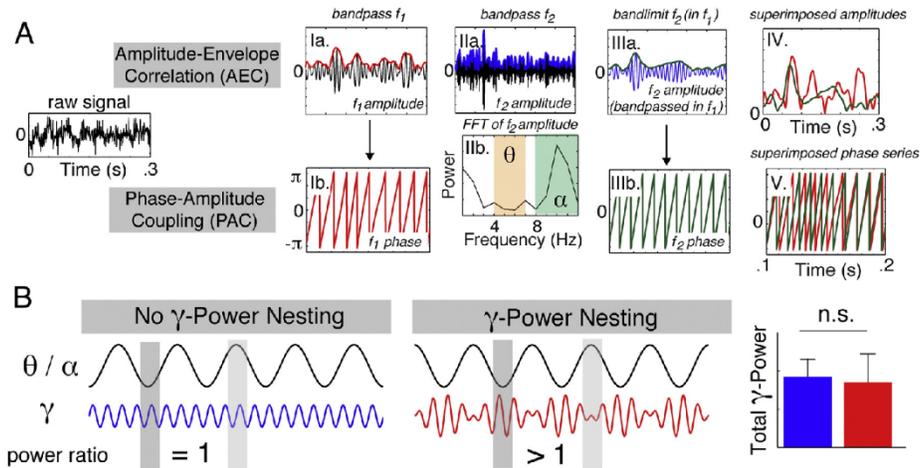


Figure 7. Analysis strategy for cross-frequency coupling analyses

(A) Raw EEG data (black) can be decomposed into different frequency bands through band-pass filtering. Instantaneous phase and amplitude estimates are extracted from the Hilbert transform and can be correlated using either linear or circular correlation analyses. (B) Visualization of the proposed coupling mechanism, where overall power remains constant, but power is modulated as a function of phase (from Helfrich et al., 2015).

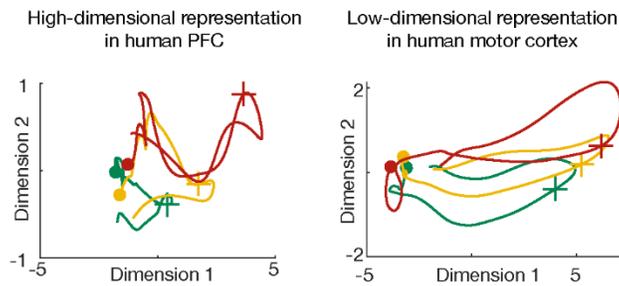


Figure 8. 2D State-space representations of cognitive dynamics

(Left) Tangled trajectories of three different conditions (varying degree of predictability) along two dimensions in human PFC indicate a complex representational pattern. (Right) In contrast, activity in motor cortex is neatly organized along a ring structure, indicating low-dimensional coding schemes (from Weber et al., 2021).

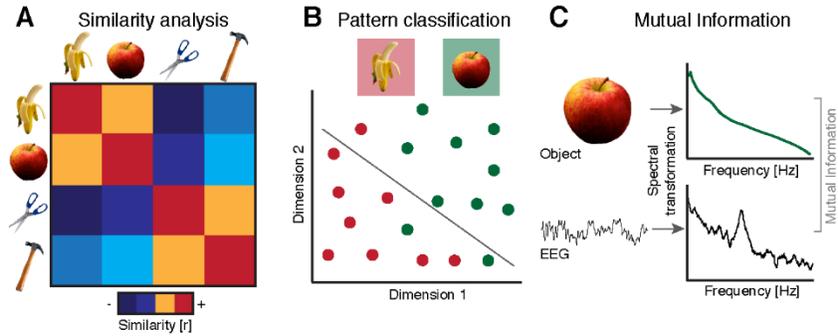


Figure 9. Information-theoretical and decoding approaches

(A) Illustration of representational similarity. Across several repetitions, similarity in neuronal representations is inferred by correlation, indicating that the same or related concepts share similar spatiotemporal patterns in electrophysiological recordings. (B) Decoding or pattern classification analysis relies on training classifiers on distinct features, which then enable a validation on a held-out testing data set. Here a linear discrimination analysis is depicted with single points representing observations that are characterized by features along two dimensions. (C) Information-theoretical analyses between object and EEG features can be conceptualized as a non-linear correlation analysis.

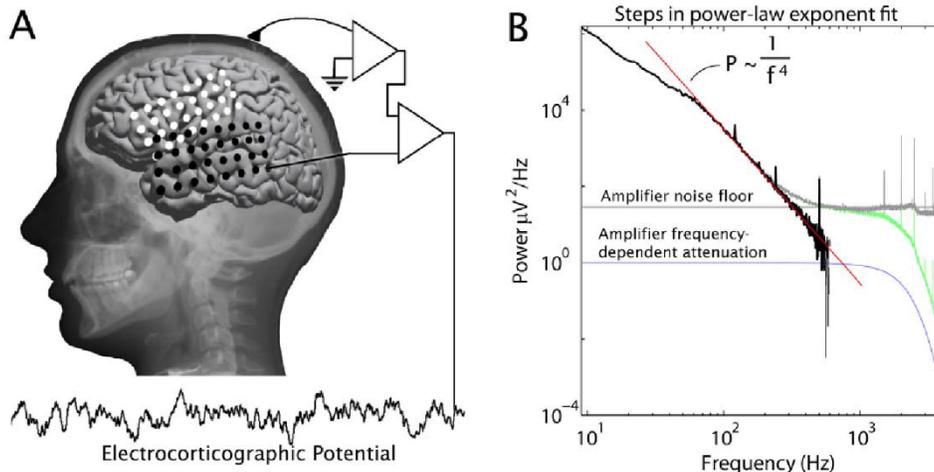


Figure 10. Power-law scaling in intracranial recordings

(A) Experimental approach. (B) $1/f^x$ power law fits of electrophysiological brain activity in relationship to recording hardware noise floors (from Miller et al., 2009).

REFERENCES

- Aru, J., Aru, J., Priesemann, V., Wibral, M., Lana, L., Pipa, G., Singer, W., & Vicente, R. (2015). Untangling cross-frequency coupling in neuroscience. *Current Opinion in Neurobiology*, *31*, 51–61.
- Baillet, S. (2017). Magnetoencephalography for brain electrophysiology and imaging. *Nature Neuroscience*, *20*(3), 327–339.
- Bastos, A. M., & Schoffelen, J.-M. (2015). A tutorial review of functional connectivity analysis methods and their interpretational pitfalls. *Frontiers in Systems Neuroscience*, *9*. doi: 10.3389/fnsys.2015.00175
- Berens, P. (2009). CircStat: A MATLAB toolbox for circular statistics. *Journal of Statistical Software*, *31*(10). doi: 10.18637/jss.v031.i10
- Berger, H. (2004). *Über das elektrenkephalogramm des menschen: Die vierzehn originalarbeiten von 1929-1938*. Berlin: Deutsche Gesellschaft für Klinische Neurophysiologie.
- Biasiucci, A., Franceschiello, B., & Murray, M. M. (2019). Electroencephalography. *Current Biology*, *29*(3), R80–R85.
- Bronstein, J. M., Tagliati, M., Alterman, R. L., Lozano, A. M., Volkmann, J., Stefani, A., ...Krack, P. (2011). Deep brain stimulation for Parkinson disease: An expert consensus and review of key issues. *Archives of Neurology*, *68*(2), 165–165.
- Buzsáki, G. (1996). The hippocampo-neocortical dialogue. *Cerebral Cortex*, *6*(2), 81–92.
- Buzsáki, G., Anastassiou, C. A., & Koch, C. (2012). The origin of extracellular fields and currents: EEG, ECoG, LFP and spikes. *Nature Reviews: Neuroscience*, *13*(6), 407–420.

- Buzsáki, G., & Draguhn, A. (2004). Neuronal oscillations in cortical networks. *Science*, *304*, 1926–1929.
- Canolty, R. T., Edwards, E., Dalal, S. S., Soltani, M., Nagarajan, S. S., Kirsch, H. E., ...Knight, R. T. (2006). High gamma power is phase-locked to theta oscillations in human neocortex. *Science*, *313*, 1626–1628.
- Canolty, R. T., & Knight, R. T. (2010). The functional role of cross-frequency coupling. *Trends in Cognitive Sciences*, *14*(11), 506–515.
- Carlson, A. A., Rutishauser, U., & Mamelak, A. N. (2018). Safety and utility of hybrid depth electrodes for seizure localization and single-unit neuronal recording. *Stereotactic and Functional Neurosurgery*, *96*(5), 311–319.
- Chaisangmongkon, W., Swaminathan, S. K., Freedman, D. J., & Wang, X.-J. (2017). Computing by robust transience: How the fronto-parietal network performs sequential, category-based decisions. *Neuron*, *93*(6), 1504–1517.
- Chari, A., Thornton, R. C., Tisdall, M. M., & Scott, R. C. (2020). Microelectrode recordings in human epilepsy: A case for clinical translation. *Brain Communications*, *2*(2). doi: 10.1093/braincomms/fcaa082
- Chini, M., Pfeffer, T., & Hanganu-Opatz, I. L. (2021). *Developmental increase of inhibition drives decorrelation of neural activity*. *bioRxiv*. doi: 10.1101/2021.07.06.451299
- Churchland, M. M., Cunningham, J. P., Kaufman, M. T., Foster, J. D., Nuyujukian, P., Ryu, S. I., & Shenoy, K. V. (2012). Neural population dynamics during reaching. *Nature*, *487*, 51–56.

- Cocchi, L., Gollo, L. L., Zalesky, A., & Breakspear, M. (2017). Criticality in the brain: A synthesis of neurobiology, models and cognition. *Progress in Neurobiology*, *158*, 132–152.
- Cohen, M. X. (2017). Where Does EEG Come from and what does it mean? *Trends in Neurosciences*, *40*(4), 208–218.
- Cole, S. R., van der Meij, R., Peterson, E. J., de Hemptinne, C., Starr, P. A., & Voytek, B. (2017). Nonsinusoidal beta oscillations reflect cortical pathophysiology in Parkinson's disease. *Journal of Neuroscience*, *37*(18):4830–4840.
- Cole, S. R., & Voytek, B. (2017). Brain oscillations and the importance of waveform shape. *Trends in Cognitive Sciences*, *21*(2), 137–149.
- Despouy, E., Curot, J., Reddy, L., Nowak, L. G., Deudon, M., Sol, J.-C., ...Barbeau, E. J. (2020). Recording local field potential and neuronal activity with tetrodes in epileptic patients. *Journal of Neuroscience Methods*, *341*. doi: 10.1016/j.jneumeth.2020.108759
- Diekelmann, S., & Born, J. (2010). The memory function of sleep. *Nature Reviews Neuroscience*, *11*(2), 114–126.
- Donoghue, T., Haller, M., Peterson, E. J., Varma, P., Sebastian, P., Gao, R., ...Voytek, B. (2020). Parameterizing neural power spectra into periodic and aperiodic components. *Nature Neuroscience*, *23*(12), 1655–1665.
- Donoghue, T., Schaworonkow, N., & Voytek, B. (2021). Methodological considerations for studying neural oscillations. *European Journal of Neuroscience*, *55*(11-12):3502-3527.
- Ebitz, R. B., & Hayden, B. Y. (2021). The population doctrine in cognitive neuroscience. *Neuron*, *109*(19), 3055–3068.

- Edwards, E., Soltani, M., Deouell, L. Y., Berger, M. S., & Knight, R. T. (2005). High gamma activity in response to deviant auditory stimuli recorded directly from human cortex. *Journal of Neurophysiology*, *94*(6), 4269–4280.
- Eichenbaum, H. (2017). Barlow versus Hebb: When is it time to abandon the notion of feature detectors and adopt the cell assembly as the unit of cognition? *Neuroscience Letters*, *680*, 88–93.
- Engel, A. K., Fries, P., & Singer, W. (2001). Dynamic predictions: Oscillations and synchrony in top-down processing. *Nature Reviews: Neuroscience*, *2*(10), 704–716.
- Engel, A. K., Gerloff, C., Hilgetag, C. C., & Nolte, G. (2013). Intrinsic coupling modes: Multiscale interactions in ongoing brain activity. *Neuron*, *80*(4), 867–886.
- Fiebelkorn, I. C., Pinsk, M. A., & Kastner, S. (2018). A dynamic interplay within the frontoparietal network underlies rhythmic spatial attention. *Neuron*, *99*(4), 842–853.
- Fisher, R., Salanova, V., Witt, T., Worth, R., Henry, T., Gross, R., ...SANTE Study Group. (2010). Electrical stimulation of the anterior nucleus of thalamus for treatment of refractory epilepsy. *Epilepsia*, *51*(5), 899–908.
- Flinker, A., Korzeniewska, A., Shestyuk, A. Y., Franaszczuk, P. J., Dronkers, N. F., Knight, R. T., & Crone, N. E. (2015). Redefining the role of Broca's area in speech. *Proceedings of the National Academy of Sciences of the United States of America*, *112*(9), 2871–2875.
- Fox, M. D., Snyder, A. Z., Vincent, J. L., Corbetta, M., Essen, D. C. V., & Raichle, M. E. (2005). The human brain is intrinsically organized into dynamic, anticorrelated functional networks. *Proceedings of the National Academy of Sciences*, *102*(27), 9673–9678.
- Fried, I., Rutishauser, U., Cerf, M., & Kreiman, G. (2014). *Single neuron studies of the human brain: probing cognition*. Cambridge, MA: MIT Press.

- Fusi, S., Miller, E. K., & Rigotti, M. (2016). Why neurons mix: High dimensionality for higher cognition. *Current Opinion in Neurobiology*, *37*, 66–74.
- Gao, R., Peterson, E. J., & Voytek, B. (2017). Inferring synaptic excitation/inhibition balance from field potentials. *NeuroImage*, *158*, 70–78.
- Gao, R., van den Brink, R. L., Pfeffer, T., & Voytek, B. (2020). Neuronal timescales are functionally dynamic and shaped by cortical microarchitecture. *ELife*, *9*. doi: 10.7554/eLife.61277
- Gerber, E. M., Sadeh, B., Ward, A., Knight, R. T., & Deouell, L. Y. (2016). Non-sinusoidal activity can produce cross-frequency coupling in cortical signals in the absence of functional interaction between neural sources. *PloS One*, *11*(12). doi: 10.1371/journal.pone.0167351
- Gervasoni, D., Lin, S.-C., Ribeiro, S., Soares, E. S., Pantoja, J., & Nicolelis, M. A. L. (2004). Global forebrain dynamics predict rat behavioral states and their transitions. *Journal of Neuroscience*, *24*(49), 11137–11147.
- Golesorkhi, M., Gomez-Pilar, J., Tumati, S., Fraser, M., & Northoff, G. (2021). Temporal hierarchy of intrinsic neural timescales converges with spatial core-periphery organization. *Communications Biology*, *4*(1). doi: 10.1038/s42003-021-01785-z
- Goudar, V., Peysakhovich, B., Freedman, D. J., Buffalo, E. A., & Wang, X.-J. (2021). *Elucidating the neural mechanisms of learning-to-learn*. *bioRxiv*, doi: 10.1101/2021.09.02.455707
- Grootswagers, T., Wardle, S. G., & Carlson, T. A. (2017). Decoding dynamic brain patterns from evoked responses: A tutorial on multivariate pattern analysis applied to time series neuroimaging data. *Journal of Cognitive Neuroscience*, *29*(4), 677–697.

- Gross, J., Baillet, S., Barnes, G. R., Henson, R. N., Hillebrand, A., Jensen, O., ... Schoffelen, J.-M. (2013). Good practice for conducting and reporting MEG research. *NeuroImage*, *65*, 349–363.
- Handy, T. C. (2005). *Event-related potentials: A methods handbook*. Cambridge, MA: MIT Press.
- Hardstone, R., Poil, S.-S., Schiavone, G., Jansen, R., Nikulin, V. V., Mansvelder, H. D., & Linkenkaer-Hansen, K. (2012). Detrended fluctuation analysis: A scale-free view on neuronal oscillations. *Frontiers in Physiology*, *3*. doi: 103389/fphys.2012.00450
- He, B. J., Zempel, J. M., Snyder, A. Z., & Raichle, M. E. (2010). The temporal structures and functional significance of scale-free brain activity. *Neuron*, *66*(3), 353–369.
- Hebart, M. N., & Baker, C. I. (2018). Deconstructing multivariate decoding for the study of brain function. *NeuroImage*, *180*, 4–18.
- Hefft, S., Brandt, A., Zwick, S., von Elverfeldt, D., Mader, I., Cordeiro, J., Trippel, M., Blumberg, J., & Schulze-Bonhage, A. (2013). Safety of hybrid electrodes for single-neuron recordings in humans. *Neurosurgery*, *73*(1), 78–85.
- Helfrich, R. F., Fiebelkorn, I. C., Szczepanski, S. M., Lin, J. J., Parvizi, J., Knight, R. T., & Kastner, S. (2018). Neural mechanisms of sustained attention are rhythmic. *Neuron*, *99*(4), 854–865..
- Helfrich, R. F., Herrmann, C. S., Engel, A. K., & Schneider, T. R. (2015). Different coupling modes mediate cortical cross-frequency interactions. *NeuroImage*, *140*, 76–82.
- Helfrich, R. F., & Knight, R. T. (2019). Cognitive neurophysiology of the prefrontal cortex. *Handbook of Clinical Neurology*, *163*, 35–59.

- Helfrich, R. F., Lendner, J. D., Mander, B. A., Guillen, H., Paff, M., Mnatsakanyan, L., ...Knight, R. T. (2019). Bidirectional prefrontal-hippocampal dynamics organize information transfer during sleep in humans. *Nature Communications*, *10*(1). doi: 10.1038/s41467-019-11444-x
- Helfrich, R. F., Mander, B. A., Jagust, W. J., Knight, R. T., & Walker, M. P. (2018). Old brains come uncoupled in sleep: Slow wave-spindle synchrony, brain atrophy, and forgetting. *Neuron*, *97*(1), 221–230.
- Hipp, J. F., Hawellek, D. J., Corbetta, M., Siegel, M., & Engel, A. K. (2012). Large-scale cortical correlation structure of spontaneous oscillatory activity. *Nature Neuroscience*, *15*(6), 884–890.
- Ince, R. A. A., Giordano, B. L., Kayser, C., Rousselet, G. A., Gross, J., & Schyns, P. G. (2017). A statistical framework for neuroimaging data analysis based on mutual information estimated via a gaussian copula. *Human Brain Mapping*, *38*(3), 1541–1573.
- Ito, T., Hearne, L. J., & Cole, M. W. (2020). A cortical hierarchy of localized and distributed processes revealed via dissociation of task activations, connectivity changes, and intrinsic timescales. *NeuroImage*, *221*. doi: 10.1016/j.neuroimage.2020.117141
- Kanth, S. T., & Ray, S. (2020). Electrocorticogram (ECoG) is highly informative in primate visual cortex. *Journal of Neuroscience*, *40*(12), 2430–2444.
- Kosciessa, J. Q., Grandy, T. H., Garrett, D. D., & Werkle-Bergner, M. (2020). Single-trial characterization of neural rhythms: Potential and challenges. *NeuroImage*, *206*. doi: 10.1016/j.neuroimage.2019.116331
- Kriegeskorte, N., Goebel, R., & Bandettini, P. (2006). Information-based functional brain mapping. *Proceedings of the National Academy of Sciences*, *103*(10), 3863–3868.

- Lachaux, J. P., Rodriguez, E., Martinerie, J., & Varela, F. J. (1999). Measuring phase synchrony in brain signals. *Human Brain Mapping*, 8(4), 194–208.
- Lendner, J. D., Helfrich, R. F., Mander, B. A., Romundstad, L., Lin, J. J., Walker, M. P., Larsson, P. G., & Knight, R. T. (2020). An electrophysiological marker of arousal level in humans. *Elife*, 9. doi: 10.7554/eLife.55092
- Leszczyński, M., Barczak, A., Kajikawa, Y., Ulbert, I., Falchier, A. Y., Tal, I., ... Schroeder, C. E. (2020). Dissociation of broadband high-frequency activity and neuronal firing in the neocortex. *Science Advances*, 6(33). doi: 10.1126/sciadv.abb0977
- Linkenkaer-Hansen, K., Nikouline, V. V., Palva, J. M., & Ilmoniemi, R. J. (2001). Long-range temporal correlations and scaling behavior in human brain oscillations. *Journal of Neuroscience*, 21(4), 1370–1377.
- Lobier, M., Siebenhühner, F., Palva, S., & Palva, J. M. (2014). Phase transfer entropy: A novel phase-based measure for directed connectivity in networks coupled by oscillatory interactions. *NeuroImage*, 85(2), 853–872.
- Lopes da Silva, F. (2013). EEG and MEG: Relevance to neuroscience. *Neuron*, 80(5), 1112–1128.
- Lundqvist, M., Herman, P., Warden, M. R., Brincat, S. L., & Miller, E. K. (2018). Gamma and beta bursts during working memory readout suggest roles in its volitional control. *Nature Communications*, 9(1). doi: 10.1038/s41467-017-02791-8
- Lundqvist, M., Rose, J., Herman, P., Brincat, S. L., Buschman, T. J., & Miller, E. K. (2016). Gamma and beta bursts underlie working memory. *Neuron*, 90(1), 152–164.
- Mante, V., Sussillo, D., Shenoy, K. V., & Newsome, W. T. (2013). Context-dependent computation by recurrent dynamics in prefrontal cortex. *Nature*, 503, 78–84.

- Maris, E., & Oostenveld, R. (2007). Nonparametric statistical testing of EEG- and MEG-data. *Journal of Neuroscience Methods*, 164(1), 177–190.
- Miller, K. J., Sorensen, L. B., Ojemann, J. G., & den Nijs, M. (2009). Power-law scaling in the brain surface electric potential. *PLoS Computational Biology*, 5(12). doi: 10.1371/journal.pcbi.1000609
- Murray, J. D., Bernacchia, A., Roy, N. A., Constantinidis, C., Romo, R., & Wang, X.-J. (2017). Stable population coding for working memory coexists with heterogeneous neural dynamics in prefrontal cortex. *Proceedings of the National Academy of Sciences*, 114(2), 394–399.
- Nolte, G., Bai, O., Wheaton, L., Mari, Z., Vorbach, S., & Hallett, M. (2004). Identifying true brain interaction from EEG data using the imaginary part of coherency. *Clinical Neurophysiology*, 115(10), 2292–2307.
- Nolte, G., Ziehe, A., Nikulin, V. V., Schlögl, A., Krämer, N., Brismar, T., & Müller, K.-R. (2008). Robustly estimating the flow direction of information in complex physical systems. *Physical Review Letters*, 100(23). doi: 10.1103/PhysRevLett.100.234101
- Palva, J. M., Zhigalov, A., Hirvonen, J., Korhonen, O., Linkenkaer-Hansen, K., & Palva, S. (2013). Neuronal long-range temporal correlations and avalanche dynamics are correlated with behavioral scaling laws. *Proceedings of the National Academy of Sciences*, 110(9), 3585–3590.
- Panzeri, S., Macke, J. H., Gross, J., & Kayser, C. (2015). Neural population coding: Combining insights from microscopic and mass signals. *Trends in Cognitive Sciences*, 19(3), 162–172.

- Parvizi, J., & Kastner, S. (2018). Promises and limitations of human intracranial electroencephalography. *Nature Neuroscience*, *21*(4), 474–483.
- Pascual-Marqui, R. D., Michel, C. M., & Lehmann, D. (1994). Low resolution electromagnetic tomography: A new method for localizing electrical activity in the brain. *International Journal of Psychophysiology*, *18*(1), 49–65.
- Piai, V., Dahlsätt, K., & Maris, E. (2015). Statistically comparing EEG/MEG waveforms through successive significant univariate tests: How bad can it be? *Psychophysiology*, *52*(3), 440–443.
- Polich, J. (2007). Updating P300: An integrative theory of P3a and P3b. *Clinical Neurophysiology*, *118*(10), 2128–2148.
- Prerau, M. J., Brown, R. E., Bianchi, M. T., Ellenbogen, J. M., & Purdon, P. L. (2017). Sleep neurophysiological dynamics through the lens of multitaper spectral analysis. *Physiology*, *32*(1), 60–92.
- Pritchard, W. S. (1992). The brain in fractal time: 1/f-like power spectrum scaling of the human electroencephalogram. *International Journal of Neuroscience*, *66*(1-2), 119–129.
- Quiñero Quiroga, R., & Panzeri, S. (2009). Extracting information from neuronal populations: Information theory and decoding approaches. *Nature Reviews. Neuroscience*, *10*(3), 173–185.
- Quinn, A. J., Lopes-dos-Santos, V., Dupret, D., Nobre, A. C., & Woolrich, M. W. (2021). EMD: Empirical mode decomposition and Hilbert-Huang spectral analyses in python. *Journal of Open Source Software*, *6*(59). doi: 10.21105/joss.02977

- Raut, R. V., Snyder, A. Z., & Raichle, M. E. (2020). Hierarchical dynamics as a macroscopic organizing principle of the human brain. *Proceedings of the National Academy of Sciences of the United States of America*, *117*(34), 20890–20897.
- Ray, S., & Maunsell, J. H. R. (2011). Different origins of gamma rhythm and high-gamma activity in macaque visual cortex. *PLoS Biology*, *9*(4). doi: 10.1371/journal.pbio.1000610
- Rich, E. L., & Wallis, J. D. (2017). Spatiotemporal dynamics of information encoding revealed in orbitofrontal high-gamma. *Nature Communications*, *8*(1). doi: 10.1038/s41467-017-01253-5
- Rigotti, M., Barak, O., Warden, M. R., Wang, X.-J., Daw, N. D., Miller, E. K., & Fusi, S. (2013). The importance of mixed selectivity in complex cognitive tasks. *Nature*, *497*, 585–590.
- Rutishauser, U., Schuman, E. M., & Mamelak, A. N. (2006). Online detection and sorting of extracellularly recorded action potentials in human medial temporal lobe recordings, in vivo. *Journal of Neuroscience Methods*, *154*(1–2), 204–224.
- Saxena, S., & Cunningham, J. P. (2019). Towards the neural population doctrine. *Current Opinion in Neurobiology*, *55*, 103–111.
- Seth, A. K., Barrett, A. B., & Barnett, L. (2015). Granger causality analysis in neuroscience and neuroimaging. *Journal of Neuroscience*, *35*(8), 3293–3297.
- Shannon, C. E., & Weaver, W. (1998). The mathematical theory of communication. Urbana, IL: University of Illinois Press.
- Shenoy, K. V., Sahani, M., & Churchland, M. M. (2013). Cortical control of arm movements: A dynamical systems perspective. *Annual Review of Neuroscience*, *36*(1), 337–359.
- Siegel, M., Donner, T. H., & Engel, A. K. (2012). Spectral fingerprints of large-scale neuronal interactions. *Nature Reviews. Neuroscience*, *13*(2), 121–134.

- Siegel, M., Warden, M. R., & Miller, E. K. (2009). Phase-dependent neuronal coding of objects in short-term memory. *Proceedings of the National Academy of Sciences of the United States of America*, *106*(50), 21341–21346.
- Staresina, B. P., Bergmann, T. O., Bonnefond, M., van der Meij, R., Jensen, O., Deuker, L., ...Fell, J. (2015). Hierarchical nesting of slow oscillations, spindles and ripples in the human hippocampus during sleep. *Nature Neuroscience*, *18*(11), 1679–1686.
- Stolk, A., Brinkman, L., Vansteensel, M. J., Aarnoutse, E., Leijten, F. S., Dijkerman, C. H., ...Toni, I. (2019). Electrocorticographic dissociation of alpha and beta rhythmic activity in the human sensorimotor system. *ELife*, *8*. doi: 10.7554/eLife.48065
- Tagliazucchi, E., von Wegner, F., Morzelewski, A., Brodbeck, V., Jahnke, K., & Laufs, H. (2013). Breakdown of long-range temporal dependence in default mode and attention networks during deep sleep. *Proceedings of the National Academy of Sciences*, *110*(38), 15419–15424.
- Tagliazucchi, E., Chialvo, D. R., Siniatchkin, M., Amico, E., Brichant, J. F., Bonhomme, V., Noirhomme, Q., Laufs, H., & Laureys, S. (2016). Large-scale signatures of unconsciousness are consistent with a departure from critical dynamics. *Journal of the Royal Society Interface*, *13*(114). doi: 10.1098/rsif.2015.1027
- Van Veen, B. D., van Drongelen, W., Yuchtman, M., & Suzuki, A. (1997). Localization of brain electrical activity via linearly constrained minimum variance spatial filtering. *IEEE Transactions on Bio-Medical Engineering*, *44*(9), 867–880.
- Varela, F., Lachaux, J. P., Rodriguez, E., & Martinerie, J. (2001). The Brainweb: Phase synchronization and large-scale integration. *Nature Reviews. Neuroscience*, *2*(4), 229–239.

- Ville, D. V. D., Britz, J., & Michel, C. M. (2010). EEG microstate sequences in healthy humans at rest reveal scale-free dynamics. *Proceedings of the National Academy of Sciences*, *107*(42), 18179–18184.
- Voytek, B., Canolty, R. T., Shestyuk, A., Crone, N. E., Parvizi, J., & Knight, R. T. (2010). Shifts in gamma phase-amplitude coupling frequency from theta to alpha over posterior cortex during visual tasks. *Frontiers in Human Neuroscience*, *4*. doi: 10.3389/fnhum.2010.00191
- Voytek, B., Kramer, M. A., Case, J., Lepage, K. Q., Tempesta, Z. R., Knight, R. T., & Gazzaley, A. (2015). Age-related changes in 1/f neural electrophysiological noise. *Journal of Neuroscience*, *35*(38), 13257–13265.
- Vyas, S., Golub, M. D., Sussillo, D., & Shenoy, K. V. (2020). Computation through neural population dynamics. *Annual Review of Neuroscience*, *43*, 249–275.
- Wainio-Theberge, S., Wolff, A., & Northoff, G. (2021). Dynamic relationships between spontaneous and evoked electrophysiological activity. *Communications Biology*, *4*(1). doi: 10.1038/s42003-021-02240-9
- Wainio-Theberge S, Wolff A, Gomez-Pilar J, Zhang J, Northoff G. (2022). Variability and task-responsiveness of electrophysiological dynamics: Scale-free stability and oscillatory flexibility. *NeuroImage*, *256*. doi: 10.1016/j.neuroimage.2022.119245
- Warden, M. R., & Miller, E. K. (2007). The representation of multiple objects in prefrontal neuronal delay activity. *Cerebral Cortex*, *17*(S1), 41-50.
- Warden, M. R., & Miller, E. K. (2010). Task-dependent changes in short-term memory in the prefrontal cortex. *Journal of Neuroscience*, *30*(47), 15801–15810.

- Watrous, A. J., & Buchanan, R. J. (2020). The oscillatory reconstruction algorithm adaptively identifies frequency bands to improve spectral decomposition in human and rodent neural recordings. *Journal of Neurophysiology*, *124*(6), 1914–1922.
- Weber, J., Solbakk, A.-K., Blenkmann, A. O., Llorens, A., Funderud, I., Leske, S., ... Helfrich, R. F. (2021). Population coding and oscillatory subspace synchronization integrate context into actions. *bioRxiv*. doi: 10.1101/2021.12.17.473118
- Wen, H., & Liu, Z. (2016). Separating fractal and oscillatory components in the power spectrum of neurophysiological signal. *Brain Topography*, *29*(1), 13–26.
- Wolff, A., Berberian, N., Golesorkhi, M., Gomez-Pilar, J., Zilio, F., & Northoff, G. (2022). Intrinsic neural timescales: Temporal integration and segregation. *Trends in Cognitive Sciences*, *26*(2), 159–173.
- Yuste, R. (2015). From the neuron doctrine to neural networks. *Nature Reviews. Neuroscience*, *16*(8), 487–497.
- Zhang, J, Huang, Z., Chen, Y., Zhang, J., Ghinda, D., Nikolova, Y., ...Northoff, G. (2018). Breakdown in the temporal and spatial organization of spontaneous brain activity during general anesthesia. *Human Brain Mapping*, *39*(5), 2035–2046.
- Zijlmans, M., Zweiphenning, W., & van Klink, N. (2019). Changing concepts in presurgical assessment for epilepsy surgery. *Nature Reviews. Neurology*, *15*(10), 594–606.