



HAL
open science

Statistical Frequency-Dependent Analysis of Trial-to-Trial Variability in Single Time Series by Recurrence Plots

Tamara Tošić, Kristin K. Sellers, Flavio Fröhlich, Mariia Fedotenkova, Peter Beim Graben, Axel Hutt

► **To cite this version:**

Tamara Tošić, Kristin K. Sellers, Flavio Fröhlich, Mariia Fedotenkova, Peter Beim Graben, et al.. Statistical Frequency-Dependent Analysis of Trial-to-Trial Variability in Single Time Series by Recurrence Plots. *Frontiers in Systems Neuroscience*, 2016, 9 (184), 10.3389/fnsys.2015.00184 . hal-01159664v2

HAL Id: hal-01159664

<https://inria.hal.science/hal-01159664v2>

Submitted on 27 Jan 2016

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Statistical frequency-dependent analysis of trial-to-trial variability in single time series by recurrence plots

Tamara Tošić^{1,2,3*}, Kristin K. Sellers^{4,5}, Flavio Fröhlich^{4,5,6,7,8}, Mariia Fedotenkova^{1,2,3},
Peter beim Graben^{9,10} and Axel Hutt^{1,2,3}

¹Team Neurosys, Inria, ²Loria, CNRS, UMR nō 7503 and ³Université de Lorraine, Loria, UMR nō 7503, Villers-lès-Nancy, France

⁴Department of Psychiatry, ⁵Neurobiology Curriculum, ⁶Department of Cell Biology and Physiology,

⁷Department of Biomedical Engineering and ⁸Neuroscience Center, University of North Carolina at Chapel Hill, Chapel Hill, North Carolina, USA

⁹Department of German Studies and Linguistics and ¹⁰Bernstein Center for Computational Neuroscience, Berlin, Germany

Abstract

For decades, research in neuroscience has supported the hypothesis that brain dynamics exhibits recurrent metastable states connected by transients, which together encode fundamental neural information processing. To understand the system's dynamics it is important to detect such recurrence domains, but it is challenging to extract them from experimental neuroscience datasets due to the large trial-to-trial variability. The proposed methodology extracts recurrent metastable states in univariate time series by transforming datasets into their time-frequency representations and computing recurrence plots based on instantaneous spectral power values in various frequency bands. Additionally, a new statistical inference analysis compares different trial recurrence plots with corresponding surrogates to obtain statistically significant recurrent structures. This combination of methods is validated by applying it to two artificial datasets. In a final study of visually-evoked Local Field Potentials in partially anesthetized ferrets, the methodology is able to reveal recurrence structures of neural responses with trial-to-trial variability. Focusing on different frequency bands, the δ -band activity is much less recurrent than α -band activity. Moreover, α -activity is susceptible to pre-stimuli, while δ -activity is much less sensitive to pre-stimuli. This difference in recurrence structures in different frequency bands indicates diverse underlying information processing steps in the brain.

Keywords: Trial-to-trial variability, Time-frequency analysis, Local Field Potentials, Recurrence plot analysis, Statistical inference, Surrogate data, Anaesthesia, Ferret

I. INTRODUCTION

Investigation of metastable states (MS) and transients of complex dynamical systems has become increasingly important over the last decades. In this context, dynamical systems spend longer time intervals in MSs than in transients between MSs. The large interest in studying such states comes from the belief that a complex temporal behaviour of systems may be decomposed into a simple sequence of alternating MSs and transients between them. This reduced description is a model that captures the essential dynamic elements of rather complex underlying dynamics. Applications range from spin glasses [1] to molecular configurations [2] and geoscientific applications [3]. In neuroscience, the related concept of sequential metastable attractors has received increasing attention in the last years [4]–[9]. Primarily, works are motivated by the experimental observation of signal features showing alternations of dynamical behaviour at fast and slow time scales [10]–[13].

Originally the concept of metastability refers to slow relaxation dynamics in statistical physics [1], [14]. In a much wider sense, this notion is nowadays used for regions in the phase space of a dynamical system

with relatively large dwell that are connected by transients [4], [5], [9]. Paradigmatic examples for those MSs are *almost invariant sets* [15] and *recurrence domains* [16], such as saddles connected by heteroclinic trajectories [17] or the “wings” of the Lorenz attractor [18]. For this attractor in particular, it is attractive itself and has two recurrence domains centered around two unstable foci. Geometrically, these domains are spatially separated and the system’s trajectory alternately approaches to and departs from the foci. The system spends much longer time in the vicinity of a focus compared to transient intervals between the two foci. Therefore, one may refer to a Lorenz wing as to a MS: the system remains for a longer time in one partition cell of the phase space before it performs a rapid transition to another partition cell of the phase space. A MS is thus identified with a recurrence domain, while non-recurrent portions of a trajectory can be compared with transients.

In neuroscience, metastability assumed increasing experimental evidence over recent years. [19], [20] observed sequences of metastable electroencephalogram (EEG) topographies, which they called *brain microstates*. [10], [11], [21] argued that components of the event-related brain potentials (ERPs) reflecting perceptual and cognitive processes could be identified with metastable brain states. Mazor and Laurent, for instance, reported sequences of metastable states in a reconstructed activation space of the locust’s neural odor circuit [12]. [13] were able to detect metastable states in epileptic EEG time series through spectral clustering methods, and most recently, [6] revealed metastable transition networks in the recovery from anesthesia. Consequently, to understand underlying neural mechanisms much better, it is necessary to develop advanced techniques to detect these recurrence structures in experimental time series.

For the identification of metastability in time series, their characteristic slow time scales must be separated from the fast dynamics of phase space trajectories. The method known as *Perron clustering* [2], separates the system’s phase states into partitions that can approximate Markov chain states [1], [2], [13], [15], [22]. Applying spectral clustering methods to the resulting transition matrix yields the time scales of the process, while their corresponding (left-)eigenvectors allow the unification of cells into a partition of metastable states [13], [22]. Another approach by [10] utilizes the slowing-down of the system’s trajectory in the vicinity of saddles by means of phase space clustering. Most recently, beim Graben and Hutt suggested to combine recurrence plot techniques and symbolic dynamics in order to partition a system’s phase space into its recurrence domains [16], [21]. The application of the latter method to experimental event-related potentials has identified metastable attractors to so-called ERP-components, known to reflect cognitive processing stages in neural information processing.

Developing novel analysis tools for representation and tracking of nonlinear transient patterns faces numerous challenges, such as reducing the signal dimensionality while preserving the information significant for the detection task or building methods robust to acquisition noise. Recurrence analysis has been used for identifying transient patterns in experimental EEG [23], for classifying patients based on EEG time series [24] and for prediction of responses during anaesthesia [25]. A key feature of recurrence analysis is to identify sequential states in a multi-dimensional signal space, as shown in most previous studies [16], [21]. If the experimental data under study is multi-dimensional, for instance a multi-channel EEG recording, the data serves directly as the input to the recurrence analysis. However, it is not valid to compute recurrence plots in the case of univariate time series and hence the data can not be analysed directly. Therefore, it is necessary to transform the univariate signal to a multivariate (multi-dimensional) signal. Typically this is done by delay-embedding techniques [26], [27] inspired by Takens’ theorem [28]. The corresponding embedding dimension and delay time in these techniques are chosen rather independent from the dynamic features of the data since typically these are not known *a priori*.

In neuroscience, patterns occurring in certain frequency bands play distinct roles in neural information processing [29], [30]. We argue that this additional knowledge can be taken into account and the present work proposes a novel technique based on time-frequency representations of univariate signals. Here, the signal is transformed into its time-frequency representation of spectral power which spans a new phase space in which the signal trajectory evolves. Hence, one may call this transformation *spectral power embedding* since the new phase space encodes instantaneous power in certain frequency bands. The additional advantage of this approach is that it permits to analyse the recurrence structure of data in

selected frequency bands. For completeness, we mention that a signal is fully defined by its instantaneous amplitude and phase.

In this work we propose a new method for the detection of metastable states in univariate neural signals. To obtain statistically significant evidence of recurrence structures in signals, we conduct a statistical test over the set of novel, frequency-selective recurrence plots (RP). Below we describe methodologies for building such frequency-selective RPs and performing statistical inference tests. These tests indicate how stable the recurrence plots are with respect to trial-to-trial variability. This novel statistical evaluation is necessary in the analysis of neurophysiological data, since trial-to-trial variability is a well-known experimental finding in such signals. In our work we analyse synthetic transient oscillations and one state variable of the Lorenz attractor involving acquisition noise to validate the methodology. Finally, the study of experimental Local Field Potentials obtained in partially anesthetized ferrets (*Mustela putorius furo*) during a visual stimulus experiment allows to extract new insights into neural information processing. For instance, we show that temporal recurrence occurs in the α -frequency band but not in the δ -frequency band. This result suggests that in the α -band the brain processes the information step-wise (*state by state*) while no step-wise process is performed in the δ -band.

II. MATERIAL & METHODS

In this section we introduce the novel method for studying temporal recurrences common in recurrence plots of different trials. Section II-A introduces classical recurrence plots and describes corresponding parameters. Then, we provide a novel method to compute recurrence plots from their time-frequency representations. In Sec. II-B we propose the statistical test method that analyses the similarity of RPs and finds their statistically significant parts. Finally, in Sec. II-C we describe the datasets used in this work.

A. Recurrence plots and novel time-frequency representations

Recurrence is a fundamental property of dynamical systems which characterises the behaviour of the system in phase space [31]. A recurrent signal instance is a moment in time when the trajectory returns to a neighborhood of a location in phase space it has already visited previously.

Deterministic dynamical systems are described by their trajectory. A trajectory $\underline{\mathbf{x}}(t) \in \mathbb{R}^n$, $t \in \mathbb{R}$ is sampled at times $t = i\Delta t$, $i \in \{1, 2, \dots, N\}$, where Δt is the sampling time interval and N is the total number of samples. For notation simplicity, in this paper we denote the signal sample $\underline{\mathbf{x}}(i\Delta t)$ by $\underline{\mathbf{x}}(i)$. Then, a recurrence plot (RP) is defined as the $N \times N$ matrix \mathbf{R} , whose elements $r_{i,j}$ take values $r_{i,j} = 1$ when two trajectory samples lay within the open ball $\mathcal{B}(\epsilon)$ of radius ϵ

$$r_{i,j} = \begin{cases} 1, & \text{if } d(\underline{\mathbf{x}}(i), \underline{\mathbf{x}}(j)) < \epsilon, \\ 0, & \text{otherwise,} \end{cases} \quad (1)$$

where $d(\cdot, \cdot)$ is a distance function and $i, j \in \{1, \dots, N\}$. Hence, recurrence plots are two-dimensional binary matrices obtained by distance based thresholding and its elements take values $r_{i,j} \in \{0, 1\}$. In this work, pixels in RPs will be color-coded white for values $r_{i,j} = 0$ and black otherwise.

For an arbitrary chosen ϵ value we can not guarantee that some of the significant dynamic features are not discarded by thresholding. To minimise such a thresholding error, we compute the optimal threshold value ϵ^* which maximises the symbolic entropy for a given distance function, as proposed in [16]. In more detail, under the assumption that recurrence domains are uniformly distributed for a given recurrence plot, the method constructs disjoint and transitive symbolic recurrence plot matrices from multivariate data. This method permits to identify MSs in a recurrence plot and maps each state (and the transients between the states) to a symbol. Consequently, one maps the high-dimensional dynamics of the system to a sequence of symbols. Let p_k be the probability of the occurrence of the state k , i.e., the number of the occurrences of the symbol k divided by the total number of occurrences of all symbols. Then maximising

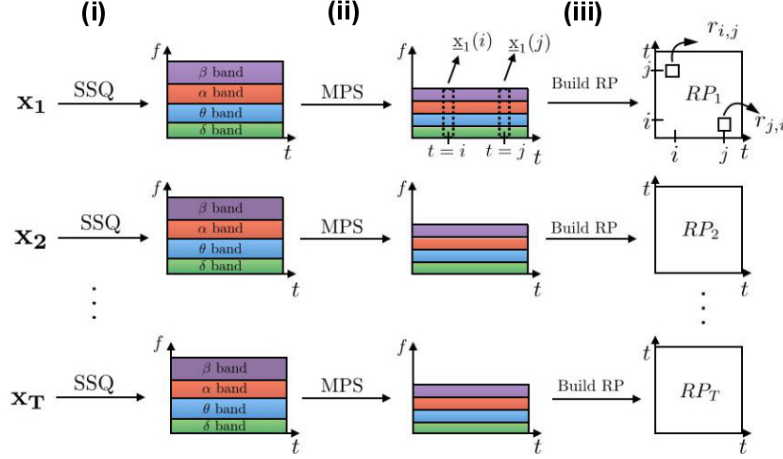


Figure 1. Building frequency-selective recurrence plots from T time series. Processing blocks are represented by arrows: (i) SSQ is a synchrosqueezing transform block used to obtain time-frequency representations of signals; (ii) the MPS block computes mean values of the power spectrum for each of the chosen frequency bands, which reduces the signal dimensionality; (iii) This signal is the basis for the recurrence analysis leading to recurrence plots $\{RP_k\}_{k=1}^T$.

the entropy

$$H(\epsilon) = -\frac{1}{S_{k,\epsilon}} \sum_{k=1}^{S_{k,\epsilon}} p_k \log(p_k), \quad (2)$$

for a range of ϵ values yields that value of ϵ for which the distribution of occurrence probabilities $\{p_k\}$ approaches uniformity, i.e., for which all states are equally probable. Here, $S_{k,\epsilon}$ is the number of states for a given ϵ . Then the optimal value

$$\epsilon^* = \arg \max_{\epsilon} H(\epsilon)$$

maximizes the entropy of the extracted symbolic sequence and hence the recurrence structure of the data. This optimal value is computed for each dataset separately.

After defining conventional RPs and computation of the optimal parameter ϵ , the remaining part of this section focusses on how to build frequency-selective recurrence plots. Many biophysiological signals have characteristic frequency signatures. For example, the human heart beats about sixty times per minute in average, i.e., at the frequency of 1Hz. Another example are eye blinks that induce signal changes in the α -frequency band (frequencies in the interval 8Hz-12Hz) in EEG recordings. To take into account the distinct signatures of spectral bands present in neural signals, we propose a novel concept for building recurrence plots from time-frequency signal representations, instead of building them directly from univariate data or constructing them by employing delay-embedding techniques. Such representations, in general adapted for nonstationary signal analysis, give insights into frequency bands of importance and provide additional flexibility to recurrence plot analysis that is not present in time-domain, for e.g., the possibility to weight the importance of some frequency bands. In the literature there are several ways to choose values of the frequency bands. We use the following frequency interval definitions: the δ -frequency band denotes the interval [0.5Hz; 4Hz], the θ -frequency band the interval [4Hz; 8Hz], α -band [8Hz; 12Hz], β -band [12Hz; 20Hz] and the γ -band denotes the interval [20Hz; 40Hz].

We build novel recurrence plots in three steps, as shown in Fig. 1: (i) we expand the set of T univariate trials $\{\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_T\}$ to their corresponding time-frequency domains; (ii) we compute the mean power of the spectrum over certain sets of frequencies. These mean power time series may be called $\mathbf{y}_j(t) \in \mathbb{R}^S$, $j = 1, \dots, T$, where S is the number of frequency bands; (iii) we compute recurrence plots by

computing distances between vectors $\mathbf{y}(i)$ and $\mathbf{y}(j)$, $i, j \in \{1, \dots, N\}$ as in Eq. (1). In the following paragraphs we describe these blocks in more detail.

The first block in Fig. 1 provides a time-frequency representation of the signal. In classical spectrogram calculations, the stronger (weaker) is the localisation of signals in time, the larger (smaller) are their localisation windows in frequency. This effect is called the uncertainty principle implied in the Fourier transform. A synchrosqueezing (SSQ) transform overcomes this deficiency by performing wavelet-based filtering and signal power reassignment to the appropriate frequencies. In addition, [32], [33] show that SSQ is superior for processing neural signals when compared to conventional spectral analysis methods, such as continuous wavelet transform or spectrogram. Hence, we use the SSQ transform defined in Section II-A.1 as the processing block (i) in Fig. 1.

The second block in Fig. 1 computes the mean value of the power spectrum (MPS) for sets of frequencies, see Section II-A.2 for details. This is one of the basic features for studying neural signals. We assume that the dynamics of the neural system encoded in frequencies is proportional to the power spectrum in sets of frequencies. This analysis step provides multi-variate time series whose dimension is equal to the resulting vector of averaged frequency bands.

Finally, in the third processing block in the figure we compute recurrence plots from the obtained time-frequency dataset as in Eq. (1). If we do not explicitly mention otherwise, we use features from all the frequency bands to compute recurrence plots. In the experimental ferret dataset, we additionally present cases when recurrence plots are calculated from the single frequency band features such as δ - or α -frequency bands, since these bands play an important role in the loss of consciousness under anesthesia.

To summarize, the proposed method for building recurrence plots from time-frequency representations grasps band-related features and allows flexibility in the analysis of particular frequency bands, which is not possible in the classical RP analysis. Our approach however requires additional computations of the synchrosqueezing transform and mean power of the spectrum.

1) *Synchrosqueezing transform*: For completeness of this work, in this section we provide the mathematical definition of the synchrosqueezing transform [32], [33], that we use as a processing block in the proposed algorithm, see Fig. 1. We presume that input signals are composed of several components with time-varying oscillatory characteristics. In other words, we assume that signals $f(t)$ can be well approximated with K signal components, $f(t) = \sum_{k=1}^K f_k(t) + e(t)$, $f_k(t) = A_k(t)e^{2\pi i\phi_k(t)}$, where $A_k(t)$ and $\phi'_k(t) = \frac{1}{2\pi} \frac{d\phi_k(t)}{dt}$ denote the amplitude and the instantaneous frequency (IF) of each component and $e(t)$ represents a small error. We assume that the components f_k have slowly time-varying amplitudes $A_k(t)$ and sufficiently smooth IFs. These conditions assure that signal components are well separated in frequencies and the complete definition is available in [34], Def. II.1 (codes available online in [35]).

Let a wavelet $\psi(t)$ be a square integrable and normalised function. Then, its scaled and time-shifted variants $\psi(\frac{t-b}{a})$ represent a set of scaled bandpass filters. In the following, we denote the frequency of one signal component by $\omega_k \approx 2\pi \frac{d\phi_k(t)}{dt}$. A Continuous wavelet transform (CWT) of the function f at scale a and time shift b is defined by $W_f(a, b) = \frac{1}{\sqrt{a}} \int_{-\infty}^{\infty} f(t) \overline{\psi(\frac{t-b}{a})} dt$, which represents a convolution of scaled and band-passed filters with the signal. The shifts of wavelet function are driven by the scale value a . For example, for the first signal component with frequency ω_1 , the value of the wavelet coefficient $W_f(a_1, b)$ spreads around the scale factor $a_1 = \frac{\omega_\psi}{\omega_1}$, where ω_ψ is the central wavelet frequency. Therefore, the estimated IF in the neighbourhood of this value of the scale is equal to the frequency ω_1 . The synchrosqueezing transform $T(\omega_q, b)$ uses estimates of the instantaneous frequency $\omega_f(a, b)$, computed for each scale-time pair (a, b) by $\omega_f(a, b) = -iW_f(a, b) \frac{\partial W_f(a, b)}{\partial b}$ to reallocate the energy of the wavelet coefficients. Let Δa_p ($\Delta\omega$) denote resolution steps in scale (frequency). Then, this transform, defined by $T(\omega_q, b) = \sum_{a_p: |\omega_f(a_p, b) - \omega_q| \leq \Delta\omega/2} W_f(a_p, b) a^{-3/2} \Delta a_p$ enhances frequency localisation of oscillating components of the signal and provides more precise time-frequency representations of the signal. In analogy to the spectrogram used in classical short-time Fourier analysis, we plot values

$$S(\omega_q, b) = |T(\omega_q, b)|^2 \quad (3)$$

for each pair (ω_q, b) in time-frequency plots, see Figs. 5(a,e) and 6(a,c).

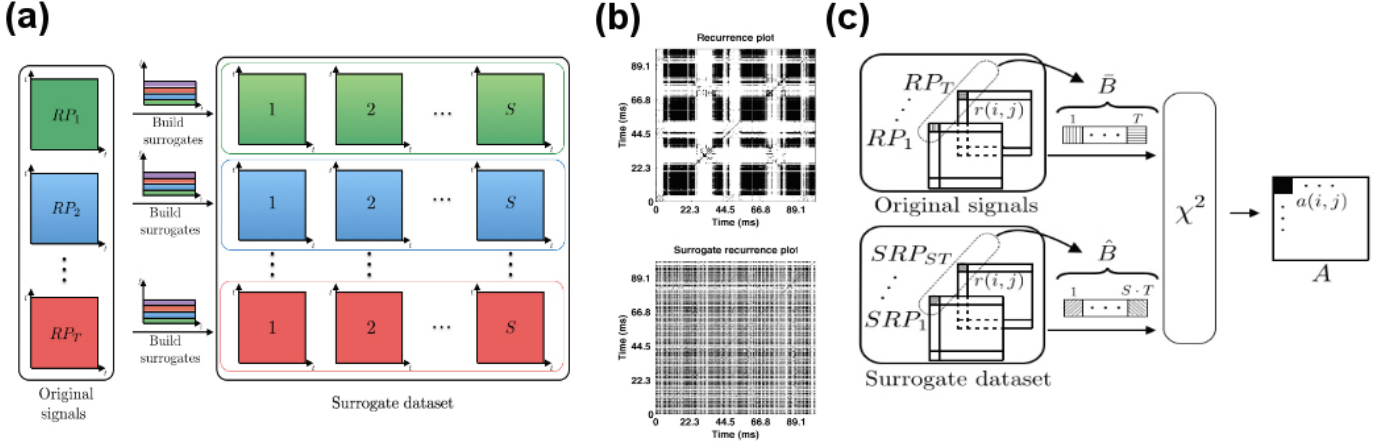


Figure 2. (a) Surrogate set construction: for each of T time-frequency signal representations used for computing the original set of RPs, one computes S surrogate RPs. The resulting S surrogate recurrence plots have the same energy in time-frequency domain. In total, there are $S \cdot T$ surrogates. (b) Examples of the RP (top) and the surrogate RP (bottom) for the transient oscillations dataset. (c) Illustration of the pixel-wise χ^2 statistical test, more details in the text.

2) *Mean power spectrum:* For each frequency band with Q components, the mean power spectrum value is defined by

$$MPS(t) = \frac{1}{Q} \sum_{q=1}^Q S(\omega_q, t), \quad (4)$$

where $S(\omega_q, t)$ is defined in Eq. (3), ω_q are frequencies of one frequency band and t is time.

B. Statistical analysis

We study statistical properties of frequency-selective RPs obtained from time-frequency trial representations. By virtue of noise effects and an expected trial-to-trial variability, recurrence plot structures are expected to vary from trial to trial. To evaluate the recurrence plots statistically, we perform a statistical inference analysis based on a classical chi-squared test [36]. To this end, we construct surrogate recurrence plots and employ an inference test.

Classically, surrogate sets of univariate signals [37] preserve some of the important features of the original time series, for example the spectrum magnitude, while they replace the phase values by a random sequence of values. The reasoning behind this randomisation is that time domain reshaping destroys non-stationarities, so the local spectral components will vary while the global spectrum remains the same. As a consequence, the mean and variance of the signal do not change [38], [39].

In this work, we build the surrogate dataset with the same power spectrum as in the original data, where the information component encoded in time is randomized, cf. Fig. 2(a). For each time index of the signal we randomly select a novel index value, such that all the index values are chosen exactly once (permutations without repetition). Then, we rearrange the time-frequency representation of trials according to the chosen index values and compute recurrence plots of surrogates by repeating steps (ii) and (iii) shown in Fig. 1. This procedure is repeated S times per trial. Figure 2(a) illustrates how to obtain the surrogate set from T trials. Examples of an original RP and a corresponding surrogate RP are provided in Fig. 2(b).

We compare pixel-related statistical measures between the set of the original recurrence plots from different trials and their surrogates to determine whether original RPs preserve the common underlying

signal dynamics in statistically significant way. This comparison is illustrated in Fig. 2 (c). In detail, we denote the set of T recurrence plots obtained from the original trial data by $\{RP_k\}_{k=1}^T$ and its surrogate set by $\{SRP_k\}_{k=1}^{S \cdot T}$. In our simulations, there are $T = 10$ trials in total, where the number of surrogates generated per trial is $S = 100$. The full set of surrogates counts $S \cdot T = 1000$ surrogate RPs. At first, we perform *pixel-wise* statistical analysis tests between the corresponding pixels of the original and the surrogate recurrence plots. Let $\bar{B} = \{r_{i,j}^{\{k\}}\}_{k=1}^T$ be the vector that consists of the set of pixels with same coordinates in the original RPs and $\hat{B} = \{r_{i,j}^{\{k\}}\}_{k=1}^{S \cdot T}$ is the corresponding vector of pixel values for surrogates. Vectors \bar{B} and \hat{B} consist of values from the set $\{0, 1\}$, since RP elements $r_{i,j}$ by definition take binary values, cf. Eq. (1). To perform a chi-square test for categorical data, we build a two-by-two contingency table. For explanation, this tables first row takes values from the original RPs and the second row contains values from surrogate RP. The first table column marks the number of values $r_{i,j} = 1$ and the second column the number of elements $r_{i,j} = 0$. The elements of this table (two rows and two columns) have the coordinates (l, m) , $l, m \in \{1, 2\}$. Then, the chi-square statistics for the pixel (i, j) , $i, j \in \{1, \dots, N\}$ is computed by

$$\chi^2(i, j) = \sum_{l \in \{1, 2\}} \sum_{m \in \{1, 2\}} \frac{(f_o^{(i,j)}(l, m) - f_e^{(i,j)}(l, m))^2}{f_e^{(i,j)}(l, m)}.$$

Here, $f_o^{(i,j)}(l, m)$ is the observed table value at the coordinate (l, m) for the pixel (i, j) and $f_e^{(i,j)}(l, m)$ is its expected frequency. The latter value is computed as $f_e^{(i,j)}(l, m) = n_r(l)n_c(m)/q$, where $n_r(l)$ is the total number of elements in the row l , $n_c(m)$ is the number of elements in the column m and q is the total number of elements in the two-by-two table. The calculated chi-square value is compared with the result in the chi-square table for predefined values of the degree of freedom $df = 1$ and the significance level $\alpha_s = 0.05$. If the calculated chi-square value is larger than the value in the table, the hypothesis that signals share the same distribution is rejected, see [36] for more details. In this work, the outcomes of chi-square tests are visually represented as matrices $A = [a_{i,j}]$, $i, j \in \{1, \dots, N\}$ whose elements take values

$$a_{i,j} = \begin{cases} 1, & \text{if distributions of trial and surrogate sets are different,} \\ 0, & \text{otherwise.} \end{cases} \quad (5)$$

In this work all the figures follow the same colour code as for illustrating recurrence plots, i.e., white pixels denote values $a_{i,j} = 0$ and black pixels stand for $a_{i,j} = 1$.

Since single elements in RPs are correlated to neighbouring elements caused by the underlying dynamics, the underlying assumption of independent recurrence matrix elements does not hold and corrections of the significance test should be applied, such as the Bonferroni correction. To this end, in the examples of artificial datasets, we have performed a t-test which is based on the hypothesis that original and surrogate signals have the same distribution of mean values. The statistics are computed based on the pixels and their mean values in a 5×5 neighbourhood around each pixel. In addition, we have applied a Bonferonni correction.

C. Datasets

To illustrate different analysis steps and to validate the power of the proposed method, we first apply the proposed algorithm to two artificial datasets. Then, the methodology is applied to experimental datasets. Single trials of these datasets are illustrated in Fig. 3 and their origin is described in detail below. For both artificial datasets, we model the tria-to-trial variability by a temporal shift of the data in time combined with additive measurement noise.

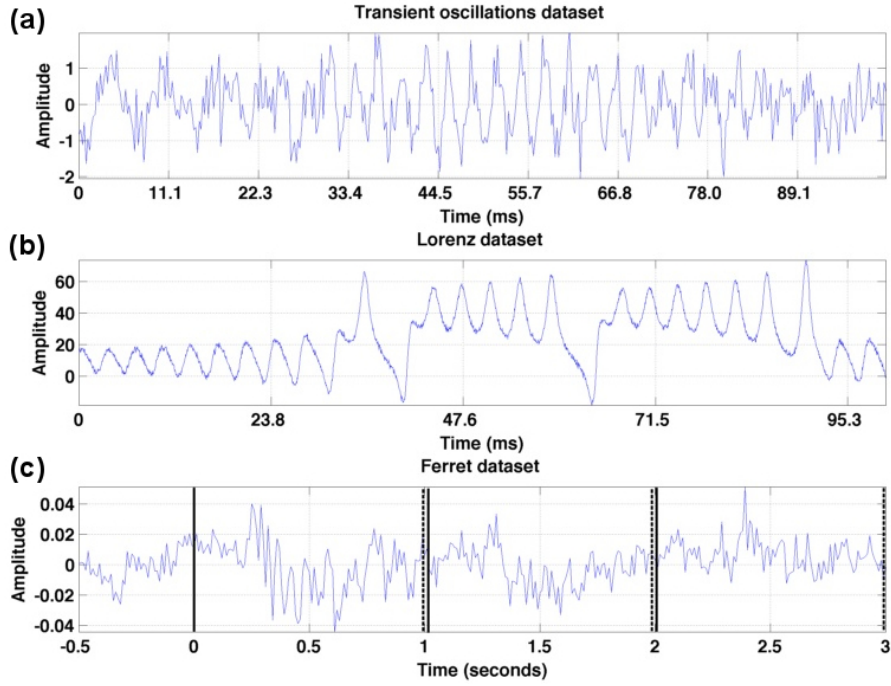


Figure 3. Illustration of the first trial in each of the three time series under study. (a) Transient oscillations (b) Lorenz attractor (c) Ferret dataset: trial from the set *session one* is recorded at a granular layer electrode. The vertical solid lines denote the stimulus onsets and the set of dashed lines mark the stimulus offsets.

1) *Transient oscillations*: A modified Lotka-Volterra model with $n = 3$ interactive elements [5], [40]

$$\frac{dx_i(t)}{dt} = x_i(t) \left(\sigma_i - \sum_{j=1}^n \rho_{i,j} x_j(t) \right), \quad (6)$$

serves as an abstract model of event-related brain potentials [21]. Here $x_i(t) \geq 0$, $i \in \{1, 2, 3\}$ is the activity rate of the element i , σ_n is the growth rate of the n -th population and ρ_i define interactions between elements. In our setup, $\sigma_1 = 1$, $\sigma_2 = 1.2$ and $\sigma_3 = 1.6$, $\rho_{ii} = 1$, $\rho_{12} = 1.33$, $\rho_{13} = 1.125$, $\rho_{21} = 0.7$, $\rho_{23} = 1.25$, $\rho_{31} = 2.1$, and $\rho_{32} = 0.83$.

The output signal $s(t)$ is a linear superposition of transient oscillations with frequencies $\nu_1 = 170\text{Hz}$, $\nu_2 = 20\text{Hz}$, $\nu_3 = 170\text{Hz}$ where at one time instance, only one of these three components is dominant, see more details below. We point out that these frequencies are chosen rather arbitrarily for an optimal illustration. The activity rate x_i defines the amplitude of the component a_i with frequency ν_i and the output signal obeys

$$s(t) = \sum_{i=1}^3 a_i(t) \sin(2\pi\nu_i t) + \xi(t) \quad , \quad a_i(t) = e^{-(x_i - \sigma_i)^2 / 2\eta_i^2}, \quad (7)$$

with $\eta_1 = 0.5$, $\eta_2 = 0.33$, $\eta_3 = 0.4$. By this construction, the amplitudes a_i increase and decrease in a certain time window outside of which they almost vanish. These windows of the three oscillation modes $i = 1, 2, 3$ do not overlap and the transitions between them are rather rapid. The variable $\xi(t)$ represents measurement noise and its random values are i.i.d. Gaussian noise with zero mean and variance 0.5. The sampling rate is 450Hz. We generate 10 trials which are time-jittered by shifting the trials by 1 sample to later instances, while each trial is subject to additive noise different in each trial. A single trial is given in Fig. 3(a).

2) *Lorenz dataset*: The Lorenz system [18] is a well-studied three-dimensional differential equation system

$$\frac{dx}{dt} = -\sigma x + \sigma y, \quad \frac{dy}{dt} = \rho x - y - xz, \quad \frac{dz}{dt} = -\beta z + xy$$

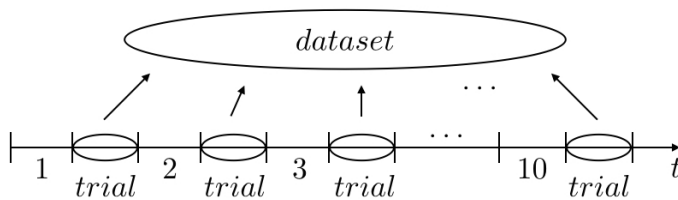


Figure 4. Experimental paradigm of the ferret experiment. In the original recording protocol [43]–[45], 10 visual stimuli are interleaved by several types of different stimuli, represented by numbers in the figure. Our dataset consists of the responses to sine-wave luminance gratings only. Other types of stimuli are the black screen (intervals marked by $\{4, 6, 9\}$), checkerboard noise for $\{2, 7\}$ or fox images stimuli. The latter set consists of the weakly spatially filtered image of foxes in the intervals marked by $\{3, 5\}$ and strongly spatially filtered image of foxes in the interval set $\{1, 5, 10\}$.

with $\sigma = 10$, $\rho = 28$, $\beta = 8/3$. Its solutions show non-trivial transient dynamics and their wings represent metastable states as explained above, cf. [16]. We study the univariate time series $x(t)$, which is the solution of the above given system of equations. This time series may represent a macroscopic measured signal such as EEG recording [41], [42], capturing activity from different metastable sources. The sampling rate is equal to 2100Hz. We generate ten signal trials time-jittered by shifting the signal by 1 sample to later instances and add i.i.d. zero mean Gaussian noise with unity variance to the signals. One trial signal is illustrated in Fig. 3(b).

3) *Ferret dataset:* The experimental dataset under study in the present work are Local Field Potential (LFP) measurements collected as described in [43]–[45]. Briefly, female ferrets were anesthetized, intubated, and underwent surgery to gain access to primary visual cortex (V1, $\sim 3\text{mm}$ anterior to lambda and 9mm lateral to the midline). Anesthesia induction was achieved with an intramuscular injection of ketamine (30mg/kg) and xylazine (1 – 2mg/kg), and anesthesia maintenance was achieved with 1.0% isoflurane (10 – 11cc, 50bpm, 100% medical grade oxygen), with continuous IV infusion of xylazine (1.5 mg/kg/hr xylazine with 4.25 mL/hr 5% dextrose lactated ringer’s). Animals were head-fixed in front of the presentation screen and a 32-channel depth probe was acutely inserted into cortex (50 microns contact spacing along the z-axis, NeuroNexus, Ann Arbor, MI) and was positioned to cover all cortical layers. The reference electrode was located on the same shank (0.5mm above the top recording site) and was positioned in 4% agar in saline above the brain. The full-field visual stimulus was presented on a 52×29 cm monitor with 120Hz refresh rate and full high-definition resolution (1920×1080 pixels, $GD235\text{Hz}$, Acer Inc, New Taipei City, Taiwan) at 47 cm distance from the animal. Each trial was 30 seconds long and consisted of three parts: (i) recording interval $[0 - 10)\text{s}$ is a baseline (screen is black); (ii) $t_s \in [10 - 20)\text{s}$ is the presentation of the sine-wave luminance gratings; (iii) $[20 - 30)\text{s}$ is ‘post-baseline’ (screen is again black). Visual stimuli were interleaved with other types of stimuli (all in randomized order), for instance with a black screen or with a strongly spatially filtered image of foxes (foxes are natural enemies of ferrets), see Fig. 4. The sine-wave luminance grating was presented at a rate of 1Hz for 10 seconds (during each 1s period, progressive frames transitioned from black to white to black and all the screen pixels had the same color for any given frame). In the subsequent analysis, we consider a subset of recordings. The dataset under study starts 0.5s before stimulus onset and lasts until 3s at the end of the third stimulus cycle. This stimulus is a black screen at $t = 0\text{s}$, $t = 2\text{s}$ and $t = 3\text{s}$ with luminance maxima at $t = 0.5\text{s}$, $t = 1.5\text{s}$ and $t = 2.5\text{s}$. A single trial is illustrated in Fig. 3(c).

Electrophysiological recordings were conducted during stimulus presentation. Unfiltered signals were first amplified with MPA8I head-stages with gain 10 (Multichannel Systems, Reutlingen, Germany), then further amplified with gain 500 (Model 3500, A-M Systems, Carlsborg, WA), digitized at 20kHz (Power 1401, Cambridge Electronic Design, Cambridge, UK), downsampled to 1kHz afterwards, and digitally stored using Spike2 software (Cambridge Electronic Design). In total, 20 trials across two sessions conducted on different days were analysed, the session sets of 10 trials are called *session one* and *session*

two in the following. Datasets are downsampled to the sampling rate equal to 100Hz.

All procedures were approved by the University of North Carolina-Chapel Hill Institutional Animal Care and Use Committee (UNC-CH IACUC) and exceed guidelines set forth by the National Institutes of Health and U.S. Department of Agriculture.

III. RESULTS

In this section, we first apply the proposed method to two artificial datasets to verify if it reveals dynamics given the noisy set of trials. After validation of the method, we apply it on the experimental dataset (ferret dataset) and study whether it well extracts the dynamics from the recorded trials.

To understand the results, we first shortly describe recurrence plots for several simple test signals. As previously mentioned, black pixels denote recurrence events and white ones its absence. All recurrence plots have a black diagonal line, by definition (see Eq. 1). Signals without any recurrence have a white square RP with a black diagonal line. Random noise signals have a random distribution of black pixels in the plot, with the exception of the black diagonal line. A simple periodic signal has a recurrence plot that consists of the black diagonal line and other black lines that are parallel to the diagonal, where the distance between them will reveal the period of the signal. More complex signals that have recurrent states may show different structures in RPs, for example, checkerboard-like patterns. These black colored fields, to which we refer as to recurrence domains, may have different sizes and shapes. For two artificial datasets we expect to observe repetitive black patterns that correspond to repetitive states within signal components. For the experimental dataset, we expect to observe recurrence patterns that are directed by the onset of the visual stimulus.

A. Artificial datasets

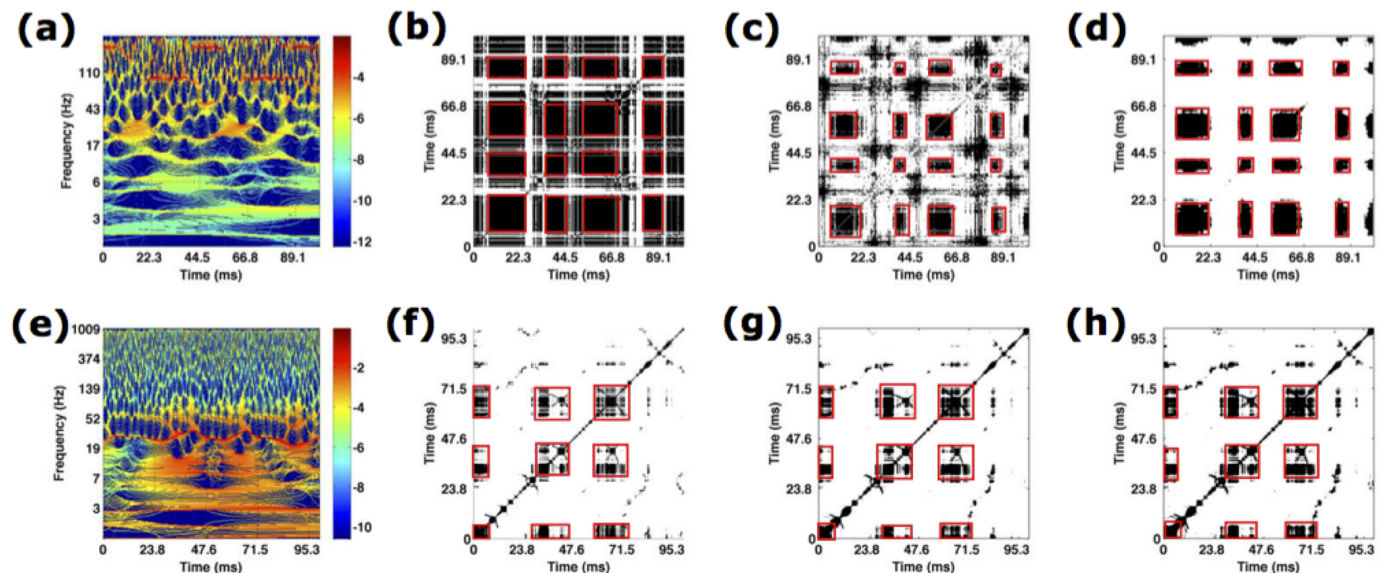


Figure 5. Illustration of time-frequency (a,e), optimal recurrence plot (b,f) for the trials of two artificial datasets shown in Fig. 3(a,b), statistically significant recurrence plots (c,g) based on 10 trials and classical t-test results (d,h). Values plotted in black represent areas where the original and the surrogate set differ significantly (rejected test). (a-d) transient oscillations; (e-h) Lorenz attractor. The red boxes mark recurrent metastable states.

We demonstrate our methodology in Fig. 5, which shows the analysis steps for the examples of transient oscillations (Fig. 5(a-d)) and the Lorenz attractor (Fig. 5(e-h)).

The time-frequency representation of one transient oscillation trial is shown in Fig. 5(a). As previously mentioned in Sec. II-C.1, the corresponding signal exhibits three periodic components. We visually inspect the figure and observe high power spectrum values around the following time windows: (i) $\{(0, 11), (42, 54), (85, 97)\}$ ms around $\nu_1 = 170$ Hz (dark red horizontal line segments); (ii) $\{(19, 26), (64, 72)\}$ ms around $\nu_2 = 20$ Hz (broad orange areas); (iii) $\{(24, 38), (74, 87)\}$ ms around $\nu_3 = 75$ Hz (dark red horizontal line segments). Note that for other trials these values may fluctuate because the frequency and time window of the current active component vary due to noise that models trial-to-trial variability. Figure 5(b) shows recurrent blocks (in black) in a single trial which correspond well to the dynamics observed in the data, cf. Fig. 5(a) and Fig. 3(a). For explanation, these recurrence blocks correspond to MSs and the white parts represent transients between them.

The time-frequency representation of one Lorenz attractor trial is given in Fig. 5(e). The approximate time intervals during which the system stays in each of the two wings are visually inspected from the power spectrum values. For the wing in time intervals $\{(0, 30), (90, 100)\}$ ms, Fig. 5(e) shows a peak at ~ 30 Hz corresponding to the oscillation frequency in the Lorenz wing, see Fig. 3(b). The other wing is reached in the time intervals $\{(40, 60), (65, 80)\}$ ms in accordance to the power peak at about 40Hz. Note that for other trials time intervals may be different due to varying trials in the set. Figure 5(f) shows recurrent blocks in a single trajectory. The recurrence blocks repeat in the correct time windows and represent the different wings, i.e., the MSs.

Time-frequency representations of single trials are the basis for the recurrence analysis leading to recurrence plots given in Fig. 5(b,f) for the respective datasets. These plots show the metastable dynamics of the transient oscillations and the Lorenz trajectories in the corresponding time windows as recurrent structures. The recurrent, i.e., repetitive, structure is visible in the illustrated trial of the corresponding data. Now, considering several trials these recurrent structures may vary due to the trial-to-trial variability. Nevertheless, to study the recurrent structure *common* to all trials, we employ the statistical inference method and extract statistically significant areas of recurrence plots, as shown in Fig. 5(c,d,g,h). The recurrent structure is obvious in these plots, reflecting the underlying recurrence structure in the artificial signals. In addition, these results demonstrate that the methodology extracts recurrence structures common in several trials, although the recurrent structure is less obvious in single trials, Fig. 5(b,c). Figure 5(d,h) show the multiple comparison-test results for both artificial datasets. The white area increases and the black areas are more focussed on the red squares, i.e. spurious recurrences (black dots) are removed and separated well from transient (white areas). Hence the multiple-comparison test improves the statistical inference.

We point out again, that the extraction of the recurrent structure from the univariate data shown in Fig. 5 is possible only by the spectral power embedding, i.e., the transformation of the univariate data into multivariate data. The subsequent preliminary statistical inference allows to identify the recurrent MSs which are common in all trials with a confidence of 0.95. By virtue of the spectral power embedding, the method permits to select certain frequency bands to study recurrence structure in specific frequency bands. This new element renders the spectral power embedding more flexible and hence superior to previous embedding techniques, such as the delay embedding based on Takens theorem. To illustrate this, the subsequent section shows results from experimental data in different frequency bands.

B. Experimental data

After studying artificially generated trials and verifying that the proposed method extracts well the dynamics features given by repetitive black structures, we now investigate whether such structures can be found in the experimental data as well.

Figure 6 provides the time-frequency representations of two single trials of the same session (a,c) and the corresponding recurrence plots (b,d). We observe a high trial-to-trial variability between both trials. This can be observed both in the time-frequency representations and the resulting recurrence plots. For instance, panel (b) shows a single recurrent state in the data except in the time window during the first

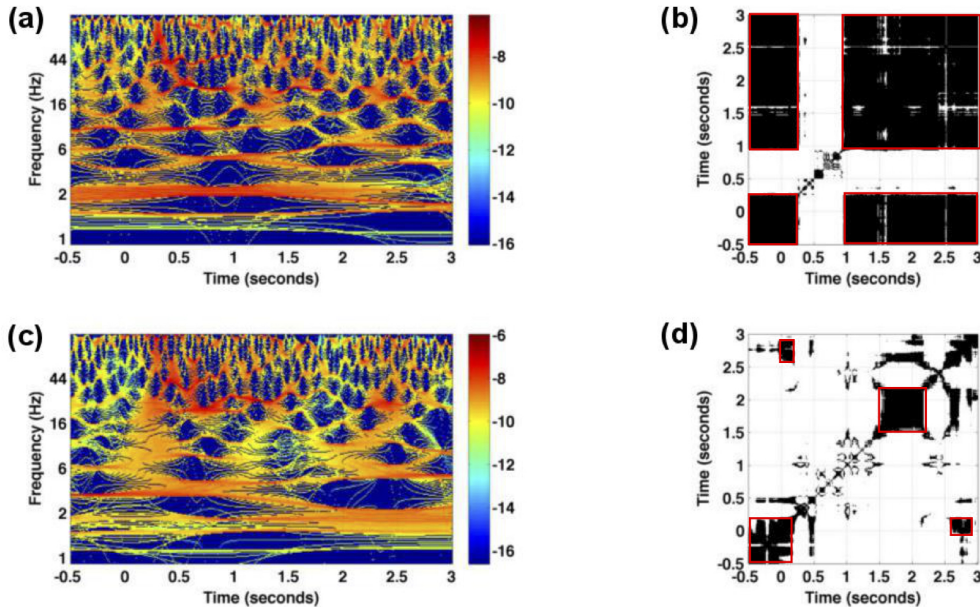


Figure 6. Illustration of trial variability for the experimental dataset *session one*. (a,c) The logarithm values of the power spectrum for two trials measured at an electrode in the granular layer; (b,d) the corresponding recurrence plots. Non-zero values of recurrence plots are encoded in black. Red boxes denote recurrent metastable states.

stimulus at $t \in [0.25\text{s}; 1\text{s}]$. Hence the system remains close to the resting state ($t < 0$) during the first stimulus. Conversely, Figure 6(d) reveals that the baseline activity, i.e., activity before stimulation in the time interval $[-0.5\text{s}; 0\text{s}]$, recurs in the interval $[2.5\text{s}; 3\text{s}]$. In addition, the activity at about $t = 2\text{s}$ resembles the activity just after $t = 2.5\text{s}$. These different findings for two trials are surprising since the experimental presentation of the visual stimulus is well-controlled and the stimulus is simple enough to expect almost identical neural responses.

To reveal the recurrent structure that is common in all trials, we now study the trial-to-trial variability of recurrence plots and aim to reveal whether the signal trials preserve the same dynamical behaviour, cf. Fig. 7. Applying the statistical method, we investigate the similarities of the results obtained from ten trials measured by a single granular sensor and from the set of ten averaged signal trials, where the average is taken over 8 granular layer sensors. This analysis is done for both ferret datasets. Moreover, we detail the analysis considering particular frequency bands which are of interest for anaesthesia. To this end, we compute recurrence plots using the values of the power spectrum coefficients in the corresponding frequency bands as illustrated in Fig. 1.

Figure 7 shows the statistically significant parts of the recurrence plots for the δ - and the α -frequency band and for all frequencies (chi-square and t-test results). The figure reveals that there is no statistically significant recurrent structure in the δ frequency band in signals under study. Conversely, the α -frequency band exhibits significant recurrent structures in the single granular electrode in both datasets, cf. Fig. 7(a,b). For instance, in *session one* the first response to the stimulus at $t = 0\text{s}$ returns at $t = 1\text{s}$. Results for all frequency bands differ to results obtained in the α frequency band. The differences are dependent on the experimental sessions suggesting the presence of strong recurrences in bands different to α and δ or strong noise artefacts. To gain further insights into the dependence on frequency bands, we consider single trials which represent spatial averages of time series from adjacent granular layers. This average denoises the time series. Figure 7(c,d) shows the corresponding results. Figure 7(c,d) show results from data in both experimental sessions revealing a similar recurrence structure now. Considering all frequency bands yields

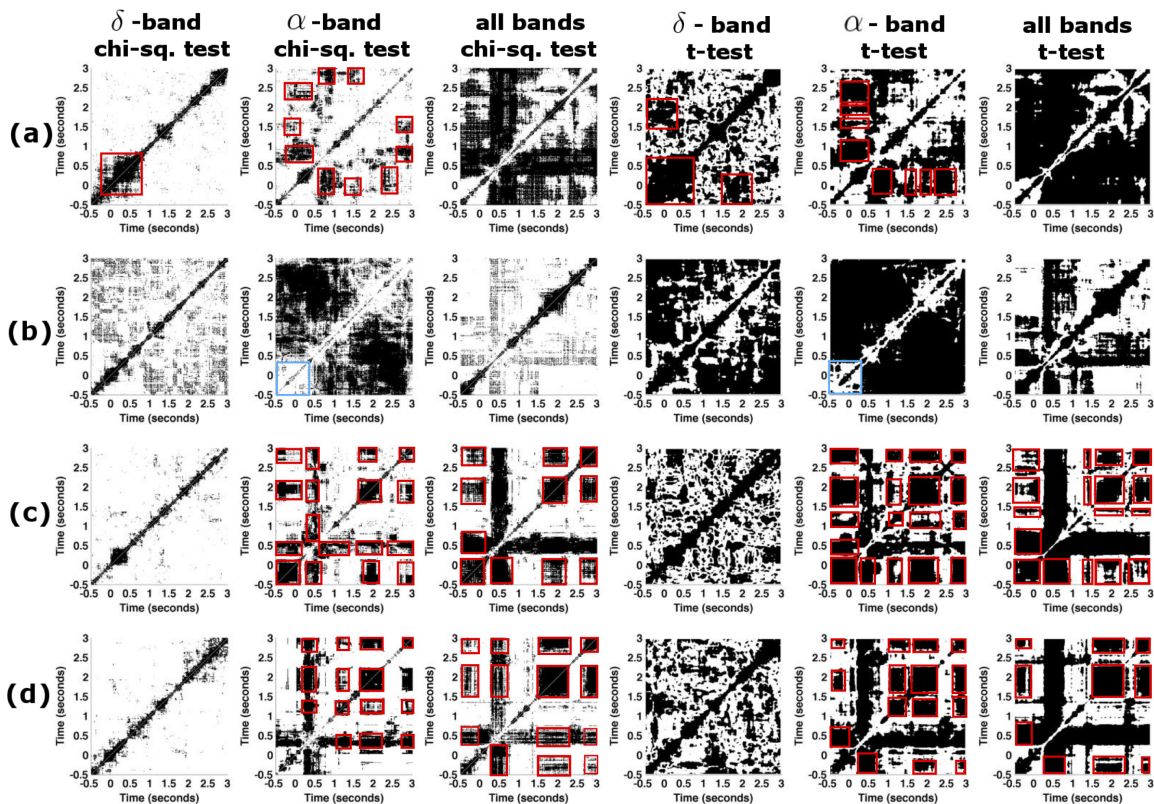


Figure 7. Statistically significant parts of recurrence plots for different ferret datasets. (a) Results from a dataset of *session one* that consists of 10 trial recordings of the single granular sensor s_{10} . (b) Same as in (a) for the dataset *session two*. (c) Results are calculated based on 10 averaged trials in dataset *session one*, where each trial is averaged over 8 granular layer time series. (d) Same as in (c) for dataset *session two*. Red boxes indicate MSs, blue boxes in (b) indicate a prominent lack of recurrence in the α -band.

recurrences similar to the one obtained in the α -band. Specifically, the prominent cross-shaped structure located at $t = 0.5$ s indicates a MS common to all data with $t \geq 0.5$ s. Additional recurrences occur in the time intervals $[-0.5; 0]$ ms, $[1.7; 2.2]$ ms and $[2.7; 3.0]$ ms. These results are consistent in the chi-square significance tests and the t-test involving corrections for multiple comparison. At last, we mention the prominent lack of recurrence in the baseline time interval observed in a single electrode in *session one*, cf. Fig. 7(b). Since it does neither occur in *session one* nor in the spatially averaged data shown in Fig. 7(c,d), it appears to be spurious and is neglected.

The experimental paradigm includes visual presentations of stimulus types in a randomised order. The previous paragraphs show neurophysiological responses to the sine-wave stimulus only. To gain further insight into the trial-to-trial variability subject to various pre-stimuli, we have selected two subsets of sine-wave trials that have two different preceding stimuli, namely the 'black screen' subset denoted by *subset one* and the subset of 'strongly spatially filtered version of foxes' denoted by *subset two*. *Subset one* includes the trials $\{4, 6, 9\}$, while *subset two* is composed of the trials $\{1, 5, 10\}$ of datasets *session one* and *session two*. The comparison of various pre-stimuli data is done by the chi-squared difference measure based on recurrence plots of both subsets. Figure 8 shows the statistically significant recurrences that are common in stimulus responses on both types of pre-stimuli. The diagonal lines are absent from figures, which suggests that at each time instance two comparison signals differ. Poor but visible recurrent structures in δ -band are grouped into two distinct blocks which distinguish the activity before the stimuli

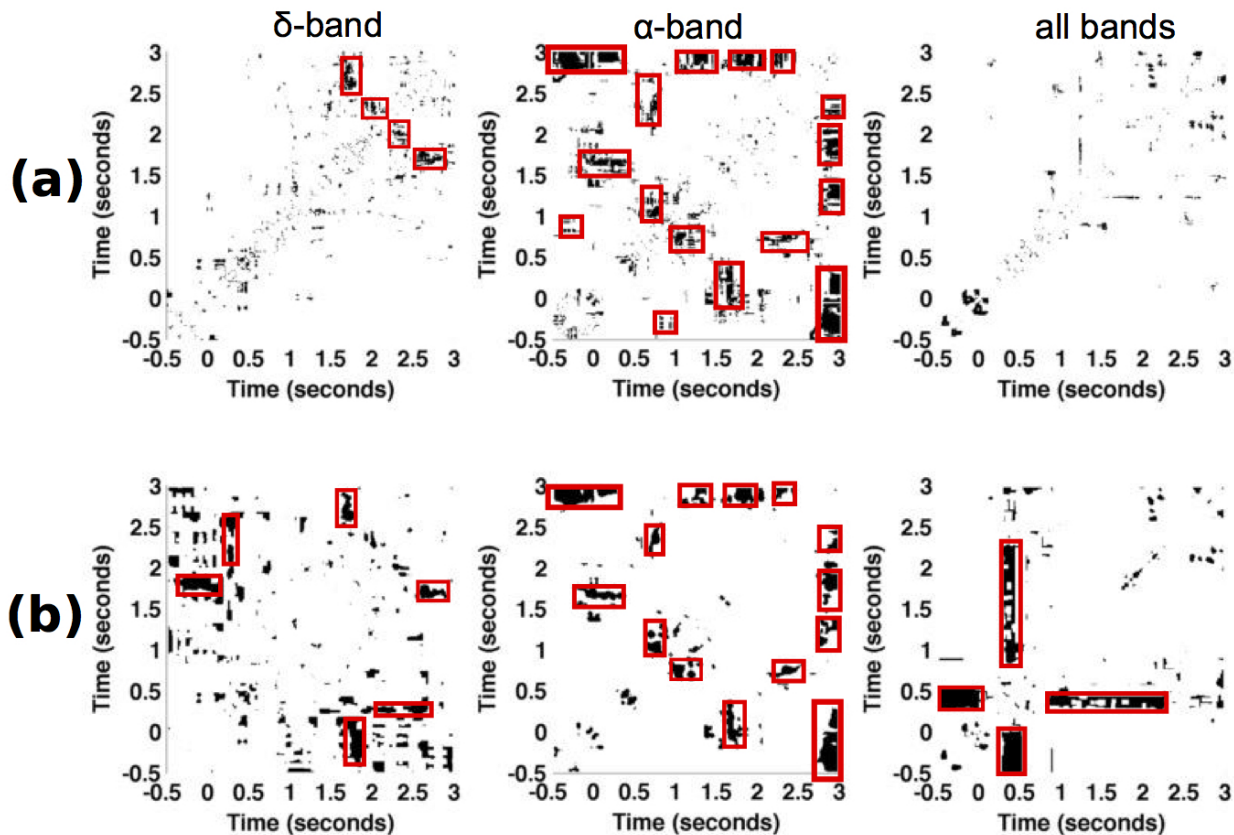


Figure 8. Illustration of the influence of different visual pre-stimuli on resulting plots in experimental dataset. Statistically significant areas of recurrence plots are obtained by (a) pixel-wise chi-square tests and (b) t-test between trials with pre-stimulus *black screen* and those whose pre-stimulus is *fox image*. The data is taken from a single granular layer electrode in datasets *session one* and *session two* together. Both the pre-stimulus *black screen* and *fox image* have occurred three times among the 10 trials in each dataset. Significantly different values are coded as white pixels, statistically similar values are coded as black pixels.

(around $t = 0$ s) and during the stimuli, for $t \in (1, 3)$ s. In α -band, the figure shows more prominent recurrences, such as the patterns around $t = 1.5$ s and $t = 2.7$ s. We point out that recurrences within δ - and α -bands do not overlap, except in the pre-stimuli period, for $t \in \{-0.5, 0\}$ s. Finally, considering all the frequency bands together does not reveal significant similarities of two pre-stimuli. Results from the chi-square test and the t-test involving multiple comparison correction are similar. However, it is interesting to note that the t-test reveals more significant common recurrences than obtained with the naive chi-square test.

IV. DISCUSSION

The present work introduces a new recurrence analysis methodology for univariate time series. The first new element is the transformation to a time-frequency representation leading to a multivariate time series of spectral power. This new technique generates a new high-dimensional phase space in which the instantaneous power of the signal evolves. This high-dimensional phase space is mandatory to apply recurrence analysis. In addition, it permits to compute recurrence plots for specific frequency bands.

The second new element is the statistical analysis of recurrence plots that takes into account spurious recurrence structures and allows to suppress them. The combination of the two methods permits to extract temporal recurrence structures in data which may reflect underlying transient dynamics in a certain range of frequencies that would have been hidden in conventional methods. To our best knowledge these two techniques have not been considered before.

The first results for two artificial datasets illustrate the methodology and indicate that method detects recurrences in a variable dataset (noise-induced trial variability) by the statistical analysis as seen in Fig. 5. These results on artificial datasets prove that the method reveals underlying recurrences in a set of trials if they are present in these trials.

The subsequent analysis of single Local Field Potentials measured experimentally in ferret visual cortex reveals a high trial-to-trial variability, cf. Fig. 6. The trial-to-trial variability is surprising due to the well-controlled experiment revealing an intrinsic ongoing activity [46]. This result demonstrates that it is mandatory to take into account recurrence variability in several trials. This is done by the methodology proposed. Detailed recurrence analysis of specific frequency bands in Fig. 7 reveals missing recurrences in the δ band whereas α -activity exhibits statistically significant temporal recurrence. This important finding reflects a fundamental difference of the nature of δ - and α -activity which has been shown in previous experimental studies on the neural origin of both signal features [47]–[49]. Our results suggest that the brain may decode information processing steps in different frequency bands. This might be of importance in previous studies and may shed some new light on neural processes, such as on metastable states in EEG during the emergence from unconsciousness [6] and metastable states in bird songs [7].

The effects of pre-stimuli have been hypothesized [50], [51] and we have investigated the effect of pre-stimuli. The performed analysis is based on a rather small set of trials reflecting the responses to identical stimuli. To have sufficiently large dataset for tests, we merged trials coming from two recording sessions. We note that trials coming from two sessions may not be independent, which may introduce errors. We found negligible effects in the δ frequency bands but differences in the temporal recurrence structure in the α frequency band. This result indicates that α -activity is more sensitive to pre-stimuli than δ -activity in the experimental setup under study. This finding is in full line with previous theoretical [50] and experimental [52] studies on the importance of phase and power of prestimulus α -activity. In addition, we notice the absence of the diagonal line and other strong recurrence patterns visible in Fig. 7. This may be the result of merging trials from different sessions, which was necessary to obtain larger test set for the analysis.

The present work shows that trial-to-trial variability in neurophysiological data occurs in spite of well-controlled and simple response-driven experimental conditions and demonstrate how to extract recurrent structures nevertheless. The methodology proposed makes it necessary to choose a well-adapted technique to transform the univariate times series to a multivariate time frequency signal. In addition to our current choice of a spectral reassignment technique, we have employed a conventional wavelet technique using complex Morlet mother wavelets and performed the same recurrence analysis (results are not shown). It turns out that this conventional method does not provide high-quality extraction of transient recurrent structures, given by the reassignment method. This may result from the worse time-frequency resolution of conventional Morlet wavelets. Future work will further investigate the best choice of multi-resolution time-frequency methods. Moreover, the methodology considers surrogate data generated by a temporal random shuffling of data and hence destructing all temporal structure. Future work may include the destruction of the recurrence structure by phase randomisation in certain frequency bands [53].

To conclude, in this work we propose a novel analysis method for trial-to-trial variability of recurrence plots in univariate time series applying a novel statistical analysis technique. This extension of recurrence analysis by a statistical technique is motivated by the fact that many physiological datasets have a limited number of trials but posses the intrinsic recurrence property of patterns of interest. Inspired by the fact that particular physiological patterns very often occur in specific frequency bands, we first build novel recurrence plots from a time-frequency signal representation. A low dimensional time-frequency signal that is built by the band median filter is then used to obtain original trial recurrence plots. Next, we use a

chi-squared statistics to obtain statistically important areas of recurrence plots. The work reveals a strong trial-to-trial variability of recurrences in experimental data in spite of the well-controlled experimental paradigm. Moreover, it turns out that recurrences occur in the α -frequency band, whereas activity in the δ -frequency band does not exhibit a temporal recurrent structure indicating frequency-dependent metastable states.

V. DATA SHARING

We provide the time-series of the transient oscillation dataset and Lorenz dataset on the webpage of the corresponding author (<https://sites.google.com/site/tamtos/datasets>).

DISCLOSURE/CONFLICT-OF-INTEREST STATEMENT

The authors declare that the research was conducted in the absence of any commercial or financial relationships that could be construed as a potential conflict of interest.

AUTHOR CONTRIBUTIONS

The majority of the analysis steps and the implementations have been performed by TT. KKS and FF have provided the experimental data and neurophysiological insights. MF has contributed additional time-frequency analysis results and PbG and AH have conceived the study. All authors wrote the manuscript together.

ACKNOWLEDGEMENT

Research reported in this publication was supported by the National Institute Of Mental Health of the National Institutes of Health under Award Number R01MH101547. The content is solely the responsibility of the authors and does not necessarily represent the official views of the National Institutes of Health. AH acknowledges funding from the European Research Council for support under the European Union's Seventh Framework Programme (FP7/2007-2013)/ERC grant agreement no. 257253. PbG acknowledges support by a Heisenberg Fellowship of the German Research Foundation DFG (GR 3711/1-2).

REFERENCES

- [1] H. Larralde and F. Leyvraz, "Metastability for Markov processes with detailed balance," *Physical Review Letters*, vol. 94, 2005.
- [2] P. Deuffhard and M. Weber, "Robust Perron cluster analysis in conformation dynamics," *Linear Algebra and its Applications*, vol. 398, pp. 161 – 184, 2005.
- [3] G. Froyland, K. Padberg, M. H. England, and A. M. Treguier, "Detection of coherent oceanic structures via transfer operators," *Physical Review Letters*, vol. 98, 2007.
- [4] K. J. Friston, "Transients, metastability, and neuronal dynamics," *NeuroImage*, vol. 5, no. 2, pp. 164 – 171, 1997.
- [5] M. I. Rabinovich, R. Huerta, and G. Laurent, "Transient dynamics for neural processing," *Science*, vol. 321, no. 5885, pp. 48 – 50, 2008.
- [6] A. E. Hudson, D. P. Calderon, D. W. Pfaff, and A. Proekt, "Recovery of consciousness is mediated by a network of discrete metastable activity states," *Proceedings of the National Academy of Sciences of the U.S.A.*, vol. 111, no. 25, pp. 9283 – 9288, 2014.
- [7] I. B. Yildiz and S. J. Kiebel, "A hierarchical neuronal model for generation and online recognition of birdsongs," *PLoS Computational Biology*, vol. 7, no. 12, p. e1002303, 12 2011.
- [8] O. Oullier and J. A. S. Kelso, "Neuroeconomics and the metastable brain," *Trends in Cognitive Science*, vol. 10, no. 8, pp. 353 – 354, 2006.
- [9] E. Tognoli and J. Kelso, "The metastable brain," *Neuron*, vol. 81, no. 1, pp. 35 – 48, 2014.
- [10] A. Hutt and H. Riedel, "Analysis and modeling of quasi-stationary multivariate time series and their application to middle latency auditory evoked potentials," *Physica D*, vol. 177, no. 1-4, pp. 203 – 232, 2003.
- [11] A. Hutt, "An analytical framework for modeling evoked and event-related potentials," *International Journal of Bifurcation and Chaos*, vol. 14, no. 2, pp. 653 – 666, 2004.
- [12] O. Mazor and G. Laurent, "Transient dynamics versus fixed points in odor representations by locust antennal lobe projection neurons," *Neuron*, vol. 48, no. 4, pp. 661 – 673, 2005.
- [13] C. Allefeld, H. Atmanspacher, and J. Wackermann, "Mental states as macrostates emerging from EEG dynamics," *Chaos*, vol. 19, p. 015102, 2009.
- [14] C. G. Tokman, B. R. Hunt, and P. Wright, "Approximating invariant densities of metastable systems," *Ergodic Theory and Dynamical Systems*, vol. 31, no. 05, pp. 1345 – 1361, 2011.
- [15] G. Froyland, "Statistically optimal almost-invariant sets," *Physica D*, vol. 200, pp. 205 – 219, 2005.

- [16] P. beim Graben and A. Hutt, “Detecting recurrence domains of dynamical systems by symbolic dynamics,” *Physical Review Letters*, vol. 110, no. 15, p. 154101, 2013.
- [17] M. I. Rabinovich, R. Huerta, P. Varona, and V. S. Afraimovich, “Transient cognitive dynamics, metastability, and decision making,” *PLoS Computational Biology*, vol. 4, no. 5, p. e1000072, 2008.
- [18] E. N. Lorenz, “Deterministic nonperiodic flow,” *Journal of the Atmospheric Sciences*, vol. 20, no. 2, pp. 130 – 141, 1963.
- [19] D. Lehmann, H. Ozaki, and I. Pal, “EEG alpha map series: brain micro-states by space-oriented adaptive segmentation,” *Electroencephalography and Clinical Neurophysiology*, vol. 67, pp. 271 – 288, 1987.
- [20] J. Wackermann, D. Lehmann, C. M. Michel, and W. K. Strik, “Adaptive segmentation of spontaneous EEG map series into spatially defined microstates,” *International Journal of Psychophysiology*, vol. 14, no. 3, pp. 269 – 283, 1993.
- [21] P. beim Graben and A. Hutt, “Detecting event-related recurrences by symbolic analysis: Applications to human language processing,” *Proceedings of the Royal Society London*, vol. A373, p. 20140089, 2015.
- [22] B. Gaveau and L. S. Schulman, “Multiple phases in stochastic dynamics: Geometry and probabilities,” *Physical Reviews E*, vol. 73, 2006.
- [23] R. Shalhaf, H. Behnam, J. Sleigh, D. Steyn-Ross, and M. Steyn-Ross, “Frontal-temporal synchronization of eeg signals quantified by order patterns cross recurrence analysis during propofol anesthesia,” *IEEE Trans Neural Syst Rehabil Eng*, vol. 23, no. 3, pp. 468–74, 2015.
- [24] D. McCarthy, N. Punjabi, P. Kim, C. Frilot, and A. Marino, “Recurrence analysis of the eeg during sleep accurately identifies subjects with mental health symptoms,” *Psychiatry Res.*, vol. 224, no. 3, pp. 335–40, 2014.
- [25] L. Huang, W. Wang, and S. Singare, “Recurrence quantification analysis of eeg predicts responses to incision during anesthesia,” in *Neural Information Processing*, ser. Lecture Notes in Computer Science, vol. 4234. Springer Berlin Heidelberg, 2006, pp. 58–65.
- [26] C. Webber and J. Zbilut, “Dynamical assessment of physiological systems and states using recurrence plot strategies,” *J. Appl. Physiol.*, vol. 76, pp. 965–973, 1994.
- [27] J. Iwanski and E. Bradley, “Recurrence plot analysis: to embed or not to embed ?” *Chaos*, vol. 8, pp. 861–871, 1998.
- [28] F. Takens, “Detecting strange attractors in turbulence,” in *Dynamical Systems and Turbulence*, ser. Lecture Notes in Mathematics, vol. 898. Springer Berlin Heidelberg, 1981, pp. 366–381.
- [29] E. R. Kandel, J. H. Schwartz, and T. M. Jessell, *Principles of Neural Science 4th ed.* McGraw-Hill, New York, 2000.
- [30] A. Schnitzler and J. Gross, “Normal and pathological oscillatory communication in the brain,” *Nature Reviews Neuroscience*, vol. 4, no. 6, p. 285296, 2005.
- [31] H. Poincaré, “Sur la probleme des trois corps et les équations de la dynamique,” *Acta Mathematica*, vol. 13, pp. 1–271, 1890.
- [32] S. Meignen, T. Oberlin, and S. McLaughlin, “A new algorithm for multicomponent signals analysis based on synchrosqueezing: With an application to signal sampling and denoising,” *IEEE Trans. on Signal Processing*, vol. 60, no. 11, pp. 5787–5798, 2012.
- [33] F. Auger, P. Flandrin, Y.-T. Lin, S. McLaughlin, S. Meignen, T. Oberlin, and H.-T. Wu, “Time-frequency reassignment and synchrosqueezing,” *IEEE Signal Processing Magazine*, vol. 80, pp. 32–41, 2013.
- [34] G. Thakur, E. Brevdo, N.-S. Fučkar, and H.-T. Wu, “The synchrosqueezing algorithm for time-varying spectral analysis: robustness properties and new paleoclimate applications,” *Signal Processing*, vol. 93, pp. 1079 –1094, 2013.
- [35] G. Thakur, “Synchrosqueezing transform codes,” 2013, <https://github.com/ebrevdo/synchrosqueezing>. [Online]. Available: <https://github.com/ebrevdo/synchrosqueezing>
- [36] F. Yates, “Contingency tables involving small numbers and the χ^2 test,” *Supplement to the Journal of the Royal Statistical Society*, vol. 1, no. 2, pp. 217–235, 1934.
- [37] T. Schreiber and A. Schmitz, “Surrogate time series,” *Physica D*, vol. 142, no. 11, p. 346382, 2000.
- [38] C. Richard, A. Ferrari, H. Amoud, P. Honeine, P. Flandrin, and P. Borgnat, “Statistical hypothesis testing with time-frequency surrogates to check signal stationarity,” *IEEE International Conference on Acoustics, Speech and Signal Processing*, pp. 3666–3669, 2010.
- [39] P. Borgnat, P. Flandrin, P. Honeine, C. Richard, and J. Xiao, “Testing Stationarity With Surrogates : A Time-Frequency Approach,” *IEEE Trans. on Sig. Process.*, vol. 58, no. 7, pp. 3459–3470, 2010.
- [40] M. Rabinovich, R. Huerta, P. Varona, and V. Afraimovich, “Transient cognitive dynamics, metastability and decision making,” *PLOS Computational Biology*, vol. 4, 2008.
- [41] C. Skarda and W. Freeman, “How brains make chaos in order to make sense of the world,” *Behav. Brain Sci.*, vol. 10, no. 2, p. 161, 1987.
- [42] E. Basar, *Brain Function and Oscillations*. Springer Heidelberg, 2006, ch. Chaos in Brain Function, pp. 169–198.
- [43] K. Sellers, D. V. Bennett, A. Hutt, and F. Frohlich, “Anesthesia differentially modulates spontaneous network dynamics by cortical area and layer,” *Journal of Neurophysiology*, vol. 110, no. 12, pp. 2739–2751, 2013.
- [44] K. Sellers, D. V. Bennett, and F. Frohlich, “Frequency-Band Signatures of Visual Responses to Naturalistic Input in Ferret Primary Visual Cortex during Free Viewing,” *Brain Research*, vol. 1598, pp. 31–45, 2015.
- [45] K. Sellers, D. Bennett, A. Hutt, J. Williams, and F. Frohlich, “Awake versus anesthetized: Layer-specific sensory processing in visual cortex and functional connectivity between cortical areas,” *J. Neurophysiol.*, p. in press, 2015.
- [46] A. Arieli, D. Shoham, R. Hildesheim, and A. Grinvald, “Coherent spatio-temporal pattern of on-going activity revealed by real-time optical imaging coupled with single unit recording in the cat visual cortex,” *J. Neurophysiol.*, vol. 73, pp. 2072–2093, 1995.
- [47] M. Hashemi, A. Hutt, and J. Sleigh, “Anesthetic action on extra-synaptic receptors: effects in neural population models of EEG activity,” *J. Front. Syst. Neurosci.*, vol. 8, no. 232, 2014.
- [48] M. Alkire, R. J. Haier, and J. H. Fallon, “Toward a unified theory of narcosis: brain imaging evidence for a thalamocortical switch as the neurophysiologic basis of anesthetic-induced unconsciousness,” *Conscious Cogn.*, vol. 9, pp. 370–386, 2000.
- [49] S. Ching, A. Cimenser, P. L. Purdon, E. N. Brown, and N. J. Kopell, “Thalamocortical model for a propofol-induced-rhythm associated with loss of consciousness,” *Proc. Natl. Acad. Sci. USA*, vol. 107, no. 52, pp. 22 665–22 670, 2010.
- [50] M. Lundqvist, P. Herman, and A. Lansner, “Effect of prestimulus alpha power, phase, and synchronization on stimulus detection rates in a biophysical attractor network model,” *J. Neurosci.*, vol. 33, no. 29, pp. 11 917–11 824, 2013.

- [51] R. Van Rullen, N. Busch, J. Drewes, and J. Dubois, "Ongoing eeg phase as a trial-by-trial predictor of perceptual and attentional variability," *Front. Psychol.*, vol. 2, p. 60, 2011.
- [52] V. Romei, V. Brodbeck, C. Michel, A. Amedi, A. Pascual-Leone, and G. Thut, "Spontaneous fluctuations in posterior alpha-band eeg activity reflect variability in excitability of human visual areas," *Cereb. Cortex*, vol. 18, pp. 2010–2018, 2008.
- [53] C. Li, G. Ding, G. Wu, and C. Poon, "Band-phase-randomized surrogate data reveal high-frequency chaos in heart rate variability," *Conf. Proc. IEEE Eng. Med. Biol. Soc.*, vol. 2010, pp. 2806–2809, 2010.