




Unsupervised identification of states from voltage recordings of neural networks

Davide Garolini, Andreas Vitalis, Amedeo Cafilisch*

Department of Biochemistry, University of Zurich, Winterthurerstrasse 190, CH-8057 Zurich, Switzerland



GRAPHICAL ABSTRACT



ARTICLE INFO

Keywords:

Multi-neuron simulation
Neural state identification
Network inference
GABAergic interneurons
NetSAP
SAPPHIRE plot

ABSTRACT

Background: Modern techniques for multi-neuronal recording produce large amounts of data. There is no automatic procedure for the identification of states in recurrent voltage patterns.

New method: We propose NetSAP (Network States And Pathways), a data-driven analysis method that is able to recognize multi-neuron voltage patterns (states). To capture the subtle differences between snapshots in voltage recordings, NetSAP infers the underlying functional neural network in a time-resolved manner with a sliding window approach. Then NetSAP identifies states from a reordering of the time series of inferred networks according to a user-defined metric. The procedure for unsupervised identification of states was developed originally for the analysis of molecular dynamics simulations of proteins.

Results: We tested NetSAP on neural network simulations of GABAergic inhibitory interneurons. Most simulation parameters are chosen to reproduce literature observations, and we keep noise terms as control parameters to regulate the coherence of the simulated signals. NetSAP is able to identify multiple states even in the case of high internal noise and low signal coherence. We provide evidence that NetSAP is robust for networks with up to about 50% of the neurons spiking randomly. NetSAP is scalable and its code is open source.

Comparison with existing methods: NetSAP outperforms common analysis techniques, such as PCA and k-means clustering, on a simulated recording of voltage traces of 50 neurons.

Conclusions: NetSAP analysis is an efficient tool to identify voltage patterns from neuronal recordings.

* Corresponding author.

E-mail address: caflisch@bioc.uzh.ch (A. Cafilisch).

URL: <http://www.biochem-caflisch.uzh.ch/> (A. Cafilisch).

<https://doi.org/10.1016/j.jneumeth.2019.01.019>

Received 21 August 2018; Received in revised form 31 January 2019; Accepted 31 January 2019

Available online 23 February 2019

0165-0270/ © 2019 Elsevier B.V. All rights reserved.

1. Introduction

Biological neural networks exhibit intricate wiring architectures, exceptional variability at the level of single neurons and highly specific temporal dynamics. The most-used techniques for multi-neuronal recordings are two-photon calcium imaging (Denk et al., 1990; Wang et al., 2003; Stosiek et al., 2003; Helmchen and Denk, 2005), multi-electrode arrays (Du et al., 2011; Buzsáki, 2004; Najafi et al., 1990; Henze et al., 2000), and, more recently, light-sheet fluorescence microscopy (Ahrens et al., 2013; Dunn et al., 2016). These experiments produce a very large amount of data with high temporal (more than 15 kHz multi-channel sampling (Li et al., 2016)) and spatial resolutions (up to 80,000 neurons recorded simultaneously (Ahrens et al., 2013; Dunn et al., 2016)). Despite these technical advancements, a robust understanding of neural encoding is still missing (Stevenson and Kording, 2011; Rolls and Treves, 2011; Romo and de Lafuente, 2013; Cohen and Kohn, 2011; Yuste, 2015), and an unequivocal matching of stimuli with neuronal responses is not routinely achievable (Rolls and Treves, 2011; Cohen and Kohn, 2011; Bialek et al., 1991; Engel et al., 2001; Shadlen and Newsome, 1998; Harris et al., 2003).

In practice, the most common approach is to derive firing rates from calcium or voltage recordings using spike sorting or inference (Denk et al., 1990; Wang et al., 2003; Stosiek et al., 2003; Helmchen and Denk, 2005; Du et al., 2011; Buzsáki, 2004; Najafi et al., 1990; Henze et al., 2000) and analyze them using various dimensionality reduction techniques (Cunningham and Yu, 2014) such as principal component analysis (Cohen and Kohn, 2011; Harris et al., 2003; Harvey et al., 2012; Freeman et al., 2014; Lopes-dos Santos et al., 2011; Peyrache et al., 2009), clustering algorithms (like k-means (Freeman et al., 2014)), or other correlation-based techniques (Ahrens et al., 2013; Cohen and Kohn, 2011). Among the most relevant theoretical frameworks for the analysis of neural recordings we mention here point process analysis (Kass et al., 2014), generalized linear models (GLMs) (Truccolo et al., 2005), entropy-based methods (max entropy models (Schneidman et al., 2006), transfer entropy information routing (Kirst et al., 2016; Palmigiano et al., 2017)), statistical decoders (maximum likelihood estimation (Paninski, 2004; Cocco et al., 2017), maximum *a posteriori* estimation (Ota et al., 2009), Ising decoders (Schaub and Schultz, 2012)), manifold analysis (Gallego et al., 2017), trajectory-based decomposition (Harvey et al., 2012), and graph-based inference (Pitkow and Angelaki, 2017). Additionally, hidden Markov models (HMM) have been used to explain how neural networks encode “hidden” variables, which are indirectly related to external stimuli (observed variables) (Kollmorgen and Hahnloser, 2014; Vidaurre et al., 2016).

All these methods define the levels of information content that appear to be significant for specific stimuli encoded by a given neuronal population. To do so, the raw data are usually preprocessed and translated into spikes, *i.e.*, they are transformed from an analog to a digital representation (Sengupta et al., 2014). For example, one of the most successful predictors of behavioral context are GLMs (Kass et al., 2014; Truccolo et al., 2005) which usually require a binary data representation. This approach subscribes to the fundamental idea that information is represented by spike-encoded signals (Cohen and Kohn, 2011), which is known to be maximally efficient only at the population level (Boerlin and Denève, 2011; Boerlin et al., 2013). However, it has been argued that the (analog) membrane potentials contain valuable information, *e.g.*, in their sub-threshold variations (Sengupta et al., 2014; Debanne et al., 2012). Moreover, it has been found experimentally that the spiking rates follow a Poisson-like process with little regularity at the single-neuron level while the membrane potentials have been argued to exhibit significant autocorrelation (Gentet et al., 2010; Yu and Ferster, 2010). For these reasons, here we decided to focus on raw, analog signals.

We propose NetSAP (Network States and Pathways) for the analysis of voltage recordings. NetSAP is based on the automatic procedure

called SAPPHIRE (States And Pathways Projected with High REsolution), developed originally for identifying states along molecular dynamics simulations of proteins (Blöchliger et al., 2013; Blöchliger et al., 2014; Blöchliger et al., 2015; Vitalis and Cafilisch, 2012). Protein motion and neuronal recordings share a high degree of complexity which is inherent to large and flexible molecules and a large set of neurons in a network, respectively. More precisely, neural networks and proteins (*e.g.*, enzymes) can populate multiple states with specific functional roles in the dynamical system. For example, neural networks alter their functional connectivity depending on the information content which is transmitted or transformed. Similarly, the activity of an enzyme is usually modulated by small-molecule modifiers (Baici, 2015), which influence the conformation of the enzyme upon binding. These states are often metastable, meaning that they have a life time shorter than the one of the enzyme itself and longer than fluctuations due to thermal motion. In stochastic dynamical systems this implies recurrence, as each state can be visited an arbitrarily large number of times.

We are able to routinely analyze long time series, *e.g.*, long and/or many trajectories of protein motion from molecular dynamics simulations. This is due to two mitigating factors. First, the scalability of our method is $DN \log N$ where N is the number of snapshots and D is the dimensionality of an individual snapshot. Second, there are hardly any overly sensitive parameters to tune. Both advantages together allow a seamless application to large data sets. Indeed, one of the major difficulties in approaches found in the literature is parameter fitting, which often scales poorly with the amount of data available. Our work is focused on avoiding strict impositions of any particular external factor, such as the assumption of learning or fitting a specific model, in order to differentiate neural information. Moreover, we decided to apply NetSAP on raw voltage recordings as we believe these to be the most direct and in some ways challenging form of data. Nonetheless, NetSAP is similarly applicable to other experimental data such as from functional magnetic resonance imaging (fMRI) (Huettel et al., 2004) or calcium imaging (Denk et al., 1990; Wang et al., 2003; Stosiek et al., 2003; Helmchen and Denk, 2005) as the underlying data mining pipeline is general.

The NetSAP approach proposed in this manuscript is an unsupervised method intended to capture the differences between neural messages without delving into efforts to directly translate the neural code. In other words, we try to identify neural states automatically, that is, without consulting associated behavioural tasks and outcomes. We rely on the assumption that neural packages of information are exchanged under the form of dynamic functional connectivity (Preti et al., 2017) which is a common assumption in state-driven analyses of fMRI (Allen et al., 2014; Yaesoubi et al., 2015a,b) and electroencephalography (EEG) data (Hirayama et al., 2015; Allen et al., 2018). This approach has been able to differentiate clinically relevant mental disorders (Buckley et al., 2015; Kirino et al., 2017), even with spatial and temporal limitations (Hutchison et al., 2013; Hindriks et al., 2016). As mentioned above, we use a density-based cluster estimator called SAPPHIRE (Blöchliger et al., 2013; Blöchliger et al., 2014; Blöchliger et al., 2015; Vitalis and Cafilisch, 2012), which shares some characteristics with DBSCAN (Ester et al., 1996) and HDBSCAN (Campello et al., 2013, 2015; McInnes et al., 2017). During the writing of our paper, a work was published that uses HDBSCAN to find neural states from digital spike representations (Grossberger et al., 2018).

To test our analysis method we use a controlled computer simulation. For the data generation, we adopt a network model of GABAergic inhibitory interneurons (*i.e.*, expressing γ -aminobutyric acid neurotransmitter), which are considered reliable network regulators (Tóth et al., 1997; Gulyás et al., 1990; Freund and Antal, 1988; Pelkey et al., 2017). Each neuron in the simulated network receives as input an external current that mimics excitation of pyramidal neurons. The simulated system has been tailored to emulate experimental voltage recordings from the hippocampus with important physiological features




Fig. 1. Descriptive pipeline of the neural network simulation and analysis. **Simulation.** Randomly initialized oscillatory input currents are injected into each neuron of the network mimicking an excitatory field of pyramidal neurons (left). The neural network consists of ISIs partially connected to each other. Using a wiring probability $p = 0.004$, we generate around 20% inhibitory directed connections (green) resulting in randomly connected components (dark green, middle). Simulated voltage levels are recorded every millisecond, and experimental noise is added. **Analysis.** For each pair of neurons, a window of data is taken (light green boxes) around each snapshot (blue, vertical lines), and a measure of correlation or similarity, $sim(x_i, x_j)$, is computed. Thus, every snapshot gives rise to a complete graph (the inferred network), and a separate measure of similarity between snapshots sampled at different times is used to reveal the states or clusters that are hidden in the high-dimensional data set (SAPPHIRE plot (Blöchliger et al., 2013; Blöchliger et al., 2014), right). During this process, we reduce the amount of weakly informative variables using a dimensionality reduction procedure, *i.e.*, we usually keep the first 15 principal components (PCs).

like afterhyperpolarization (Wang and Buzsáki, 1996; McKay et al., 2013), excitatory/inhibitory balance (Xue et al., 2014), and theta/gamma rhythms (Buzsáki, 2002; Belluscio et al., 2012; Fries, 2015; Headley and Paré, 2017).

2. Materials and methods

This section consists of two parts. First, we generate realistic voltage recordings using a simulated neural network. Second, we take the voltage recordings and, without preprocessing, analyze them in an unsupervised manner. The main goal is the automatic identification of states without knowing the external context and by using only raw voltages. A general pipeline is depicted in Fig. 1 and described in the following.

2.1. Simulation

The graph structure in Fig. 1 shows an example of the simulated neural network. We focused on inhibitory Wang-Buzsáki neurons (Wang and Buzsáki, 1996) because they have been shown to be fundamental for population coding (Stefanelli et al., 2016; Buzsáki, 2001)

and have a clear voltage trace with typical afterhyperpolarization (Wang and Buzsáki, 1996; McKay et al., 2013) (Fig. 2a). Moreover, it has been shown that these neurons, which represent 10–15% of the neuronal population in the hippocampus, are a major determinant of the regulation of information flow by controlling the spiking in a time-windowed manner (Tóth et al., 1997; Gulyás et al., 1990; Freund and Antal, 1988; Pelkey et al., 2017). These neurons have been reported to drive the timing of systemic excitability by means of broad innervation (Pelkey et al., 2017). In the hippocampus they are, therefore, considered fundamental regulators of the oscillatory activity across different frequencies (Buzsáki, 2002). Significantly, networks of interneurons have been suggested as one of the major players in frequent neurological disorders such as schizophrenia (Marín, 2012; Lisman et al., 2008; Lewis et al., 2005) and epilepsy (Trevelyan et al., 2007; Marx et al., 2013). In both cases, a modification at the circuitry level and persistent defects in theta/gamma waves were found to be fundamental anomalies (Lisman et al., 2008; Trevelyan et al., 2007; Marx et al., 2013).

Here, we considered a system focusing on the subset of interneurons (Pelkey et al., 2017) that are disinhibitors (Freund and Buzsáki, 1996; Blasco-Ibáñez and Freund, 1995; Bezair and Soltesz, 2013; Freund and




Fig. 2. Characteristics of the simulated data set. (a) Voltage values for one single neuron during 200 ms of simulation (black). The simulated experimental error is added as pink and white noise (blue). (b) Univariate distribution of voltage levels from one interneuron during a 20-s simulation. (c) Scatter plot of spikes showing bursts of activity mimicking the theta oscillations in the input. (d) The total input current for all the 50 neurons during one second of simulation. (e) One neuronal input during one second of simulation is decomposed into its three main components: constant ($I_{c,ib}$, green), oscillatory ($I_{w,ib}$, red), and random ($I_{n,ib}$, blue) currents. (f) Excitatory/inhibitory balance. The orange line represents the average excitatory contribution of the network in terms of current ($\sum_{i=1}^n I_{ext,i}/n$, with $n = 50$). The mean inhibitory current is shown in gray ($\sum_{i=1}^n I_{syn,i}/n$) and their difference in black. (g) Firing rates (FR) of the neural population with bin sizes of 1 and 15 ms (gray and black, respectively). Frequencies on the y-axis are indicative only for the gray curve. The gamma oscillations (in gray, 40–100 Hz) are nested in theta rhythms (in black, roughly six oscillation in one second shown).

(Gulyás, 1997), i.e., they wire and inhibit other GABAergic interneurons (Chamberland and Topolnik, 2012; Tyan et al., 2014; Gulyás et al., 1996). These neurons are called interneuron-selective interneurons (ISIs), and they represent around 20% of all CA1 neurons (Bezaire and Soltesz, 2013). Even though they are usually divided into three sub-categories, we decided to focus on the specific type I subgroup of ISIs, which are known to form bundles of 10–15 neurons and target other interneurons across a wide span of areas. We considered a network of 50 neurons, containing a minimum of 10 and a maximum of 15 ISIs (Pelkey et al., 2017; Acsády et al., 1996a; Acsády et al., 1996b). In Fig. 1 we show an example of a random graph with 14 directed edges. The expected number of directed edges can be calculated, in the case of a random graph, using $p n(n - 1)$, in which p is the pairwise wiring probability and n is the number of nodes. In practice, the network was wired with uniform probability of $p = 0.004$ in order to maintain the number of outbound connections around 20% of the number of neurons while specific network topologies were ignored. We followed this approach to avoid the addition of supervised bias in the network structure, which would have required a specific definition of neural plasticity or behavioural coding.

We modeled each interneuron i with the following single-compartment current balance equation,

$$C_i \frac{dV_i}{dt} = -I_{Na,i} - I_{K,i} - I_{leak,i} - I_{syn,i} + I_{ext,i} \quad (1)$$

$$I_{ext,i} = I_{c,i} + I_{w,i} + I_{n,i} \quad (2)$$

where C_i is the membrane capacitance while $I_{leak,i}$ and $I_{ext,i}$ are the leaky and applied electric currents (external to the network), respectively. Furthermore, $I_{Na,i}$ and $I_{K,i}$ follow the Hodgkin-Huxley formulation of the ion channels (Hodgkin and Huxley, 1952). In this implementation, kinetics and maximal conductances are modified according to Wang and Buzsáki (1996) to fit hippocampal fast-spiking neurons. These cells show brief voltage afterhyperpolarization of about 15 mV after each spike (McCormick et al., 1985) (Fig. 2a and b) and can occasionally reach high frequencies (200–400 Hz (McCormick et al., 1985)) although they usually operate in the 20–100 Hz range (gamma oscillations). In our model, variations in cell features were approximated by a small variability in the C_i (extracted from a normal distribution with mean $\mu_{C_i} = 1 \mu\text{F}/\text{cm}^2$ and standard deviation $\sigma_{C_i} = 0.1 \mu\text{F}/\text{cm}^2$).

The synaptic current I_{syn} (Eq. (3)) depends on the unitless gating variable s_i , which roughly describes the fraction of open ion channels. This variable is the only inhibitory connection between interneurons. The variable s_i is bound between 0 and 1, and its dynamics are described by a single differential equation (Eq. (4)).

$$I_{syn,i} = g_{syn} s_i (V_i - V_{syn}) \quad (3)$$

$$\frac{ds_i}{dt} = \frac{\alpha(1 - s_i)}{1 + e^{-V_{pre}/2}} - \beta s_i \quad (4)$$

The conductance $g_{syn} = 0.1 \text{ mS}/\text{cm}^2$, channel opening rate $\alpha = 12$

ms^{-1} , channel closing rate $\beta = 0.1 \text{ ms}^{-1}$, and $V_{\text{syn}} = -75 \text{ mV}$ are kept constant (Wang and Buzsáki, 1996). V_i is the membrane potential of the post-synaptic neuron while V_{pre} is the pre-synaptic voltage.

The system of differential equations (1–4) was simulated using Brian2 (Goodman and Brette, 2008; Goodman et al., 2014). The input to the network is modeled as external currents (I_{ext}) that resemble the receptive field of each interneuron from different populations of excitatory pyramidal cells. This external contribution has an important role in the formation of theta oscillations in the hippocampus and is considered the major input to interneurons (Buzsáki et al., 1983). We modeled the oscillations using oscillatory currents with a frequency of 6.4 Hz. Importantly, the model naturally gives rise to gamma oscillations (25–100 Hz) nested in theta oscillations (Fig. 2g) in accordance with *in vivo* observations (Buzsáki, 2002; Belluscio et al., 2012; Fries, 2015; Headley and Paré, 2017).

The applied current $I_{\text{ext},i}$ (Eq. (2), Fig. 2e) consists of three main components: first, a constant baseline $I_{c,i}$ of $0.6 \mu\text{A}$ (green line); second, an input process with a periodic sinusoidal signal $I_{w,i}$ (semi-amplitude of $0.6 \mu\text{A}$, red line); last, an Ornstein-Uhlenbeck process (represented by the $I_{n,i}$ contribution) that adds noise to the differential equations (blue line). A fundamental aspect of the simulations is the balance between inhibitory and excitatory currents (Xue et al., 2014). This property is supported by experimental evidence (Atallah and Scanziani, 2009) and is believed to be crucial for the homeostatic coupling needed for efficient coding and computation (Lim and Goldman, 2014, 2013). In particular, we highlight here the concept that the excitatory/inhibitory balance can be positive or negative in order to effectively convey and transform information in time in a manner that goes beyond the tight homeostatic excitatory/inhibitory balance (Denève and Machens, 2016). In other words, packages of information may be propagated in the neural network using specific alterations in the excitatory/inhibitory balance (Lim and Goldman, 2014, 2013). In Fig. 2f, we show the momentary shifts in the equilibrium (black line) as the difference between the mean of all excitatory currents (orange) vs. that of all inhibitory currents (gray).

All the simulated parameters have been set following experimental findings (Pelkey et al., 2017; Wang and Buzsáki, 1996) or an ‘edge-of-chaos’ approach (van Vreeswijk and Sompolinsky, 1996). For the latter, we intended to avoid possible divergences in the information content of the neural signal. In other words, we allow neurons to spike irregularly while we still keep trace of the underlying signal which must be recurrent. We use coherence to measure voltage variability and derive adequate parameters for the simulation:

$$C_{xx'} = \frac{S_{xx'}^2}{S_x S_{x'}} \quad (5)$$

where the cross spectral density estimate $S_{xx'}$ is divided by the single spectral densities of x and x' , S_x and $S_{x'}$, respectively. The spectral densities are calculated using Welch’s method (Welch, 1967) with segments of 200 snapshots and 90% of overlap. In detail, we divided every simulation into two equal parts and computed the coherence for each neuron, *i.e.*, x is the same neuron as x' , but in the first and second halves of the data set, respectively. In this way, we can evaluate for individual frequencies how the signal of each neuron changes during the simulation. In the case of zero coherence, we expect that the first part of the simulation is completely different from the second part. For the analysis of the simulation parameters we always used simulations of 20 s (20,000 snapshots).

Finally, experimental error is simulated using white and pink noise that is added to each point of the voltage traces. The error terms were generated using the same dispersion (standard deviation) as employed in the simulated data set (around 25 mV). In Fig. 2a, this process is shown for one neuron.

2.2. NetSAP analysis

To analyze neural time series, we have extended here the SAPPHERE algorithm developed originally for the analysis of molecular dynamics simulations (Blöchliger et al., 2013; Blöchliger et al., 2014; Blöchliger et al., 2015; Vitalis and Caflich, 2012). This method is useful to sort and classify instantaneous snapshots from dynamical processes that change continuously in time. In practice, the method’s only essential hyperparameter is a distance metric between snapshots. Briefly, the algorithm starts from an arbitrary time point, subsequently adds the snapshot that is geometrically nearest to any of the snapshots already added and thus constructs a reordered time line (called progress index) without repetition. Those points in the input data that are visited extensively and come from a structurally homogeneous region will be grouped closely. A kinetic annotation (Krivov and Karplus, 2006) highlights specific states or basins based on their temporal persistence, at least when visited recurrently. In general, the density of snapshots, which is metric-dependent, should be high in the basins (states or attractors) and low in the connecting areas (‘barriers’ or transition states). This similarity clustering is highly efficient at revealing the stochastic dynamics of very complex systems, *e.g.*, when there are many small attractors found across very long recording times. Notably, the algorithm is already scalable (*i.e.*, applicable to very large data sets) and has been used to analyze millions of time points and hundreds of variables (Bacci et al., 2015).

We used mainly the minimum spanning tree (MST) construction to create the SAPPHERE plot because it has no auxiliary parameters. When computing times were unfeasible, we used an approximation, the short spanning tree (Blöchliger et al., 2013; Vitalis and Caflich, 2012), which smoothly converges to the MST with straightforward parameter choices. In particular, the number of search attempts defines how many snapshots are compared to connect the spanning tree (Fig. S1c). A larger number of attempts gives a better approximation of the MST but the dependency of the metric of interest (NMI, see below) is very weak here.

We adopted the SAPPHERE algorithm because it can elucidate subtle temporal features and resolve rare events as it preserves kinetic information (Blöchliger et al., 2013; Blöchliger et al., 2014; Blöchliger et al., 2015; Vitalis and Caflich, 2012). In particular, the SAPPHERE analysis identifies states that contribute to the preservation of information for short periods of times, which we hypothesize to constitute the physical basis for neural encoding. We assume that neural networks exchange information in finite packages that persist during prolonged and repeated input presentations. Therefore, the approach is based on analyzing the data with the goal of highlighting the functional state of the network, *i.e.*, the precise condition of the information exchange at a certain time point. How can this functional state be represented for an instantaneous time point? Here, we use a sliding window procedure as follows. For every snapshot, data from the time window centered on this point are used to calculate a correlation measure for every pair of neuron traces and we chose the inverse of the Minkowski distance for this. These data constitute a complete correlation (or similarity) matrix between neurons for each individual snapshot. The time series of these matrices provides a dynamical picture of the functioning of the network. Fed into the SAPPHERE analysis, we employ the Euclidean distance between matrices as the (dis)similarity measure. However, to counterbalance the presence of weakly informative variables, a dimensionality reduction procedure is applied to the elements of these matrices, and only the first 15 principal components are retained. The work-flow is depicted in Fig. 1.

We termed the overall procedure NetSAP (Network States And Pathways). The method belongs to the group of time-based clustering descriptors (Liao, 2005). Our aim is to infer the landscape of accessible states from observations made at discrete and generally regular time points. Broadly speaking, NetSAP serves as an unsupervised clustering algorithm that relies on transient network states to identify the system’s

local *minima* (basins or cluster centers) and their connections. Notably, NetSAP does not require a prior definition of the number of clusters or their sizes and the window length is its main impactful parameter.

To quantify the performance of NetSAP, we compare the original input annotation (true labels) and the states identified by SAPHIRE analysis (predicted labels) using a compact descriptor of the clustering performance. As shown in Results, here the true labels derive from controlled variations in the input presentation to the network. In particular, we adopt the normalized mutual information score (NMI (Strehl and Ghosh, 2002)) to evaluate the identified states:

$$\text{NMI}(\mathbf{tc}, \mathbf{pc}) = \frac{\text{MI}(\mathbf{tc}, \mathbf{pc})}{\sqrt{H_s(\mathbf{tc})H_s(\mathbf{pc})}} \quad (6)$$

where \mathbf{tc} and \mathbf{pc} denote true and predicted labels, respectively, while $H_s(\cdot)$ is the Shannon entropy and $\text{MI}(\mathbf{tc}, \mathbf{pc})$ is the mutual information between two different partitions. In case the exploratory analysis of the SAPHIRE plot highlights fewer or more states than the true annotation, we do not impose any merging or splitting. A fundamental property of clustering comparison metrics is the so-called constant baseline property (Vinh et al., 2010). If this property is not fulfilled, then the comparison metric will be positive, even in case of labels that were generated randomly. In practice, this happens when the number of data points is scarce and the number of clusters is high. Vinh et al. suggested that NMI complies with this property if $N/n_{cluster} > 100$, where N is the number of points and $n_{cluster}$ is the number of clusters (Vinh et al., 2010). Our simulated data is above this soft limit. Nevertheless, we adopted the adjusted Rand index (ARI) (Hubert and Arabie, 1985) in the case of more than four clusters (Figs. 9c–f, 10, 11).

The SAPHIRE analysis is used to identify the most relevant barriers between states or, in other words, the best separations between basins. The identification can be done by manually selecting the barriers on the cut profile of the SAPHIRE plot or by using an automatic procedure. For this latter purpose, we focus on evaluating the overlap between points on the temporal annotation of the progress index (gray dots in Fig. 3a and b). To do so we uniformly divide the x -axis in partitions and repeat this process with different numbers of divisions, so as to have a variable number of separators in different positions of the progress index (Fig. 3a). As a standard, we divide the progress index in uniform partitions 50 times by using a variable number of separators, usually between 10 and 200. In Fig. 3a, we highlight nine uniform separators of the progress index distribution with vertical red lines. The resulting 10 divisions include groups of points that are successively divided into horizontal histograms, *i.e.*, we took the y -axis (experimental time) as the traditional binning axis (Fig. 3b, 20 bins). By measuring how well two consecutive groups of points overlap, we obtain an estimate that should be proportional to the likelihoods that they come from a homogeneous basin. We take the maximal information coefficient as the measure of overlap (Reshef et al., 2011).

In practice, we estimate distributions of points with horizontal histograms and we calculate the inverted maximal information coefficient (IMIC) between consecutive pairs of vectorized histogram counts. The metric shows low values where the points are similarly distributed, while it presents peaks where there is vertical separation, *i.e.*, there is a jump between distributions of points. The IMIC between two count vectors \mathbf{a} and \mathbf{b} can be written as

$$\text{IMIC}(\mathbf{a}, \mathbf{b}) = 1 - \max_{\mathbf{AB} < n^\alpha} \frac{\text{MI}^*(\mathbf{A}, \mathbf{B})}{\log(\min(\mathbf{A}, \mathbf{B}))} \quad (7)$$

The values of the horizontal histograms \mathbf{a} and \mathbf{b} are used as numerical vectors and the scatter plot between these values is used by MIC to infer their relationship. In particular, MIC assumes the following: if there is a functional relationship between two variables, it must be possible to define a specific 2D histogram on the scatter plot which maximizes their mutual information. In Eq. (7), (\mathbf{A}, \mathbf{B}) are bins of vectors \mathbf{a} and \mathbf{b} , respectively. To limit the search space of possible bins,




Fig. 3. Procedure for the identification of the barriers. (a) The temporal trace of an 80-s simulation of 50 neurons, reordered according to the progress index, is shown in the background (gray dots, right y -axis). We divide the progress index distribution into uniform splits with different number of divisions (red points: darker to lighter colors represent 10, 25, 50, and 100 uniform divisions). Only the splits for 10 groups are shown (vertical dashed lines). For each consecutive region, two histograms in experimental time are defined, and the IMIC between them is calculated. These values are shown as the y -axis position of the red points. Colored areas correspond to histograms highlighted in panel (b). Each group of points (one for each number of divisions) is connected using segments. The resulting lines are subsequently averaged in order to create the final IMIC curve (black). The final selection of three barriers is indicated (blue vertical lines). (b) Details for the histogram construction from panel (a). Using 20 bins, we show a histogram in experimental time within every uniform progress index division. These histograms are plotted as bar plots to the side of the corresponding divisor. We use these counts to calculate pairwise IMIC values. (c) Dependence of different scores on the number of states. Using IMIC (red) and PBM index (Pakhira et al., 2004) (black), a consensus choice of four states emerges. The NMI score (blue) is an *a posteriori* validation using the true labels, *i.e.*, the ground truth.

we used n^α where n is the number of elements in \mathbf{a} and \mathbf{b} , and α is a constant (fixed at 0.4). $\text{MI}^*(\mathbf{A}, \mathbf{B})$ is the maximum mutual information over all possible bins (\mathbf{A}, \mathbf{B}) . These calculations are made using the R package *minerva* (Albanese et al., 2013).

Next, we linearly interpolate the scores found for each group of barriers, and we average them, thus obtaining the black curves in Fig. 3a and b (and the red curves in Fig. S3a–e). We use a peak finding algorithm with a span of 2500 to identify all relevant barriers and sort these *maxima* by their associated IMIC values. Considered in succession, they give rise to the corresponding profile in Fig. 3c (red, segmented line). Because the IMIC is monotonously decreasing with growing number of states, the selection of the number of barriers to use can be ambiguous (*e.g.*, Fig. S3e). For this reason, the PBM index (Pakhira et al., 2004) was consulted (Fig. 3c) as it is robust with respect to




Fig. 4. Coherence-based analysis of the simulation parameters. (a) Two simulations that differ primarily in noise currents ($0.01 \mu\text{A}$ in blue and $0.4 \mu\text{A}$ in green) are used to highlight limiting values of coherence. Unless stated otherwise, the simulation with low noise current (blue), which has mean coherence around 0.7, is used for the subsequent analysis. The lighter lines represent each single neuron used in the simulation while the darker lines are their respective means. Vertical black lines indicate the frequency values (6, 25, 36, 45, 90 Hz from left to right) for which results are shown in panels b–f. (b–f) Individual curves reflect different frequencies (green, red, black, blue, and light blue for 6, 25, 36, 45, 90 Hz, respectively). For each panel, we ran 40 simulations of 20 seconds in which all but one parameter was kept constant. The default choices were as follows: (viz., semi-amplitude of oscillatory current $I_w = 0.6 \mu\text{A}$, $I_c = 0.6 \mu\text{A}$, maximal semi-amplitude of noise current $I_n = 0.02 \mu\text{A}$, $p = 0.004$, $\sigma_{C_m} = 0.1 \mu\text{F}/\text{cm}^2$). The parameter varied is the one on the x-axis. The population firing rate (PR) is highlighted (orange line, right y-axis). Defaults are highlighted with vertical black lines. (b) Coherence increases as a function of the oscillatory current semi-amplitude (I_w). Notably, when the constant current is higher than the oscillatory current amplitude (e.g., when $I_w = 0.3 \mu\text{A}$), the coherence of the slow frequency (6 Hz, green) recovers. For further analysis we used a fixed current semi-amplitude of $0.6 \mu\text{A}$ (vertical dashed line). (c) The injected current (I_c) reduces the coherence of the signal if too high, which is in accordance with previous findings (Wang and Buzsáki, 1996). Moreover, it is consistent with panel (b), as the coherence for 6 Hz remains constant and higher than the coherence of all the other frequencies in the case of $I_c > 2I_w$. The value of the fixed current we adopted is $0.6 \mu\text{A}$. (d) As expected, the network coherence in the high frequency regime is highly dependent on the noise inserted in the differential equations. For the maximal semi-amplitude of the noisy current (I_n), we chose $0.02 \mu\text{A}$. (e) Coherence decreases as a function of the pairwise connection probability. The number of outbound connections grows rapidly with p and at 0.02 the network is already fully connected (blue vertical line). A high number of inhibitory connections increases the probability for silent neurons, which spike rarely and thus show low coherence. In order to maintain the number of self-inhibitory interneurons (ISIs) around the physiological value of 20%, we defined the connection probability as 0.004 (black vertical line). (f) Coherence levels for different standard deviations of the membrane capacitance (σ_{C_m}).

different data sets and partitions (Vendramin et al., 2019). The final selection of barriers is shown by vertical blue lines in Fig. 3a (and by dotted vertical lines in Fig. S3a–e). Generally, the resulting choice of barriers is in good accordance with a simple visual inspection.

We require one parameter for the automatic identification of peaks, which is the aforementioned minimum number of points (span) between them (for example, in Fig. 3 the minimum number is 2500 snapshots). This is a soft requirement as it is mainly necessary to avoid the selection of multiple barriers from the same peak. Since the span requirement limits the number of peaks that can be found, we use the average peak height as a cutoff for inferring the number of states. The IMIC curve is not particularly sensitive to the other parameters, i.e., the number of bins for the histogram construction and the number of uniform splits, and we keep them fixed here. As mentioned above, we used from 10 to 200 splits 50 times, and 50 bins (only 20 bins are shown in Fig. 3b to avoid overcrowding).

To be certain of the NMI robustness against random inputs, we selected the simulation with low noise (Fig. S3a) and shuffled its progress index. The resulting NMI is close to zero (Fig. S4a). As a second test, we also shuffled the true label vector, i.e., the annotation, and obtained similarly low scores (Fig. S4b). Finally, we wondered how the choice of barriers could affect the results. To investigate this possibility, we took different simulations and picked the barriers 500 times at random. Each random pick was constrained to be distant from every other pick by at

least 2500 snapshots (the same limit used for the automatic barrier identification). Afterwards, the resulting distribution of NMI values was compared with the score calculated using the automatic procedure (Z-score, Fig. S3f–j, obtaining p -values lower than 5%, see Fig. S3f–i). These statistical tests indicate that the automatic procedure for the choice of barriers is robust and performs substantially better than a random choice.

3. Results

3.1. Analysis of simulation parameters by coherence of voltage patterns

The coherence is influenced strongly by the noise current I_n (Fig. 4a). It deteriorates to values close to 0.01 for a maximal semi-amplitude of $I_n = 0.4 \mu\text{A}$ at frequencies higher than 30 Hz. An oscillatory component larger than $0.5 \mu\text{A}$ is required for a high coherence (Fig. 4b). In particular, the oscillatory current $I_{w,i}$ regulates directly the theta frequency of the wave (around 6 Hz, Fig. 2g), while the gamma ripples are defined by the finer details of the neural network, e.g., the inhibitory connections. The other components of $I_{ext,i}$ decrease signal coherence (Fig. 4c and d). In the particular cases in which $I_c > 2I_w$, we see a drop in the coherence of the theta waves (6 Hz, green in Fig. 4b–e). Nonetheless, the low frequency regime is generally more robust toward parameter variation than the higher frequency regime




Fig. 5. Simulation of a 50-neuron network according to Eqs. (1–4). (a) Input presentation sequence. Each state represents one of four inputs along the simulation of 80 s. The thin vertical lines indicate the moment in which the input currents (I_{ext}) change state. The dashed blue line highlights the transition point shown in panel b. (b) Voltage levels of 50 neurons in a time interval around a change of state (blue dashed line).

(Fig. 4b–d). Finally, the coherence of the neural simulation is not influenced by the standard deviation of the membrane capacitance irrespective of the frequency (Fig. 4f).

The connectivity or self-connectivity is critical and must be around the physiological value of 20% (Freund and Buzsáki, 1996; Blasco-Ibáñez and Freund, 1995; Bezaire and Soltesz, 2013; Freund and Gulyás, 1997), which corresponds to a connection probability $p = 0.004$. Higher values of connectivity lead to the formation of silent interneurons, i.e., neurons that spike rarely due to the high number of inhibitory inputs. *Vice versa*, lower values allow the majority of interneurons to spike continuously with high frequencies. In this way, the overall higher excitation can potentially disrupt the theta oscillations. With a 20% flat connectivity, the neurons have an overall spike frequency of 30 Hz with peaks of 80–100 Hz in the gamma ripples (Fig. 2g).

3.2. Identification of neural states by NetSAP analysis

Using the settings described in the coherence analysis, a first benchmark simulation of a 50-neuron network was performed, and voltage values were recorded for all neurons every millisecond during 80 s. Four distinct states were defined by input currents that differed slightly in timing and intensity for each neuron of the neural network. In detail, the phase and amplitude of the oscillatory input currents were taken randomly from a normal distribution having standard deviations of 0.1 radians and 0.01 μA , respectively. We subsequently simulated the four states along windows of variable lengths (between 1 s and 10 s) and connected the resulting data sets in a single time series. No other parameter was changed during the simulation, and the network




Fig. 6. Unsupervised identification of states by NetSAP analysis. (a) Using the network inference, the four different states are successfully identified by the SAPHIRE plot with almost optimal NMI of 0.92. The progress index is annotated with kinetic information reliant on a cut-based partitioning criterion (light blue curve) and the original cluster membership (top, in spectral colors). The distances between consecutive snapshots in the progress index sequence (top, black) indicate higher values before each barrier as expected. The temporal annotation (black dots) shows the real time (vertical axis) of every snapshot as a function of position in the progress index. Notably, the temporal trace highlights that the state recognition is correct also for short input presentations, such as the one found within simulation times of 60–70 s, which corresponds to only 1 s, i.e., 1000 consecutive snapshots (highlighted in green). (b) Temporal order of true states (T) and the states predicted using NetSAP (P-N) and k-means++ (Arthur and Vassilvitskii, 2007) (P-K). The latter method, using $k = 4$, is only able to recognize the blue state partially (NMI of 0.15).

connections were kept identical, similar to the one shown in Fig. 1. We did not take into consideration a possible bias when the simulation is reinitialized because, in a simulation with low noise terms, 300 ms are the maximum time needed to stabilize the spiking patterns. The simulated input sequence is shown in Fig. 5a as a continuous line between the four different states. By simply parsing visually the voltage levels, it is difficult to determine the point in which the state changes (Fig. 5b).

We applied NetSAP to the four-state simulation ($I_n = 0.01 \mu\text{A}$, mean coherence of 0.7) using 80 snapshots (80 ms) as window size for the network inference (see 2.2). This size was chosen because it is the minimum that allows the selection of a region in time that has always at least one spike. The results are robust with respect to the window size in the range from 60 ms to 120 ms (Fig. S1d). Figs. S1b and Fig. S2d show that different network inference methods affect the resulting score and computational time needed considerably. For pairwise comparisons of neural signals we decided to use the Minkowski similarity because it shows higher NMI scores than other metrics across the tested simulations (Fig. S1b). The Minkowski similarity is the inverse of the distance $d_p(X, Y) = (\sum_{i=1}^w |x_i - y_i|^p)^{1/p}$, in which x_i and y_i are vectors of neural voltages, $p = 3$, and w is the size of the sliding window.

NetSAP is able to resolve multiple neuronal states, here, the four different patterns (Fig. 6a). If we compare the temporal order of the original states (T, Fig. 6b) and the states predicted by NetSAP (P-N,




Fig. 7. PCA and SAPHIRE without network inference fail to identify the four neural states. (a) Kernel density estimation of the first two components of the principal component (PC) analysis performed on raw voltage levels. (b) Without network inference and using only Euclidean distances between snapshots, the SAPHIRE plot does not separate the states. This is expected as a purely time point-wise approach should not be able to partition an oscillatory signal encoded as dynamical correlation in time.

Fig. 6b), we can see a clear correspondence (NMI of 0.92). Conversely, a clustering procedure on the inferred functional networks managed to achieve only a partial recognition of states. Using k-means++ (Arthur and Vassilvitskii, 2007) we obtained an NMI of only 0.15 (P-K, Fig. 6b). As additional comparisons, we performed a principal component analysis (PCA) of the whole data set and constructed a SAPHIRE plot without network inference (Fig. 7a and b, respectively). Both of these controls fail to separate the four states.

Importantly, the NetsAP algorithm is able to resolve very short excursions into different states as highlighted by the green rectangle in Fig. 6a. The overall analysis procedure required 15 min (Fig. S2) on a commodity computer running Linux (Core i7-4790 CPU with 3.60 GHz x8 and 16 GB RAM). Indeed, the approximate algorithm (the so-called short spanning tree procedure) can be highly efficient and scales comfortably to millions of data points. For the data set considered here, 5000 search attempts for the construction of the short spanning tree are enough to obtain satisfactory performance (Fig. S1c).

3.3. Robustness of NetSAP

First we investigated the robustness of NetSAP for increasing noise current. We generated 20 different simulations with identical settings but using different amplitudes for the noisy input current. For frequencies higher than 6 Hz, the coherence deteriorates significantly already at a noise current of 0.1 μA (Fig. 8a). Using NetSAP and automatically selecting the barriers as described previously, we were able to find the four relevant states and to calculate an NMI score for each simulation (Fig. 8b). For further analyses, we selected five indicative simulations with the following noise levels: 0.01, 0.04, 0.12, 0.22, and 0.30 μA (gray shaded circles). Using these data sets, we estimated the




Fig. 8. Robustness of NetsAP evaluated by 20 simulations of a 50-neuron network with a variable noisy input current. (a) Coherence as a function of input noise at selected frequencies (average and standard deviation). The overall coherence is also shown (black line). (b) NMI scores as a function of coherence. Gray circles indicate the simulations used for further testing (Figs. S1 and S2), and their noise levels are indicated. The red line is a polynomial fit of degree 4 and drawn purely as a guide to the eye.

variability of NetSAP by running the analysis 35 times for each noise level changing only the random seed for the short spanning tree construction (5000 search attempts) and the automatic selection of barriers. The NMI score shows that the NetSAP analysis correctly identifies basins up to a noise level of 0.2 μA (Fig. S1a) while the standard deviation of the NMI score does not depend on the noise level. Moreover, the principal sources of variability, *i.e.*, the number of search attempts and the automatic barrier selection, can be potentially limited by manual optimization. For example, one can improve the identification of the barriers by enhancing the number of bins on the y-axis or by using more splits of the progress index distributions.

All the predicted states have been calculated using the SAPHIRE plots, and in particular using only the temporal annotations (distribution of points on the progress index). For this reason, we tested the procedure of automatic barrier selection with randomized null models (Fig. S4). In addition, we compared the NMI score generated by the automatic selection of the barriers and the distribution of NMI scores produced by random picks of divisors. The distance of the optimal score from the distribution has been estimated using Z-scores and *p*-values (Fig. S3f–j). As expected, the automatic procedure produces significantly better results than the random selection, *e.g.*, Fig. S3a–d. Instead, when the NMI score drops below a threshold of around 0.25, the selected barriers produce similar scores in comparison to a random




Fig. 9. NetSAP analysis of a neural network simulation with 20 different states sampled multiple times along a simulation of 200 seconds. The network consists of 50 neurons modeled by Eqs. (1–4). (a) State recurrence during the time course of the simulation. The average residence time for each input is 2500 snapshots (ms), and each state is visited 3–6 times. (b) Similarly to Fig. 6, the SAPHIRE plot shows in gray colors the states from dark to light. Here, the kinetic annotations (blue and red lines) illustrate likely positions of barriers. Vertical black lines on top describe the inter-snapshot distances as the height of each segment. (c) Distribution of ARI values from randomly picked barriers (light blue) and the ARI obtained using our optimization procedure (minimum distance between peaks is 7000 snapshots). Here, the p -value is 3×10^{-5} , indicating a substantially better choice in comparison to random levels. (d–f) The temporal annotation is depicted by gray points in the background. These dots are plotted as a function of their position in the progress index (x-axis) and the experimental times (y-axis). The dark red line shows the IMIC, which was interpolated to approximate a continuous line along the progress index. The blue crosses indicate the ‘barrier candidates’ and their respective height on the IMIC curve. The final selection of divisors (dotted vertical lines) is done by ranking and choosing the highest peaks (black dots at the top). Using the first divisors with the highest IMIC values, we computed the ARI score (shown as the horizontal blue line) between the inferred and true labels. As negative controls, we picked barriers randomly from a uniform distribution 500 times, and we use the same annotation to calculate the ARI (green dots, position of the barriers on the x-axis and score on the y-axis). Red dots represent the set of those random divisors with the maximum ARI score. Random shuffling of the progress index (e) or the annotation (f) results in ARI scores close to zero. (g) The temporal sequences of the true states and the states predicted using NetSAP are shown for the five states 1–5 in (a). For visual clarity, we colored in white the remaining 15 states.

pick (Fig. S3e).

We also wanted to test if the automatic barrier selection and the resulting NMI score are dependent on the number of clusters in the analyzed system. To this end, we first artificially increased or reduced the number of states in the true label vector (annotation). In practice, we divided one state into two, thereby obtaining five states overall (Fig.

S5c,f). Alternatively, we merged two states to arrive at three final states (Fig. S5b,e). Moreover, we tried to look for four states after merging two in the true annotation (Fig. S5a,d). In these cases, the better divisors in terms of NMI are generated from the random picks (red dots in Fig. S5a–c). This is expected because the temporal distribution is, in all cases, mismatched with the true label vector. These results provide




Fig. 10. Dependence of NetSAP analysis on the number of neurons and states. (a) NetSAP is applied on 400 simulations of 100 s that share the same parameters ($I_n = 0.02 \mu\text{A}$ and average visiting time between 1000 and 3000 ms) while we varied the number of states (from 2 to 40) and the number of neurons (from 10 to 105). ARI scores are shown from lower levels (blue) to higher (light green). (b) Analysis of 11 simulations of 90 s each, each having a total population of 70 neurons divided into two subpopulations that represent different states at the same time. A subset of 10 to 60 neurons can define either a 6-state or 3-state system. As an example, the first tick on the x-axis (20) corresponds to the simulation with 20 and 50 neurons receiving as input the 6-state and 3-state system, respectively.

evidence that the automatic identification makes specific use of the information about the kinetic partitioning of the states, which is desired. The best results using random splits are superior because, in essence, they are optimized using the real annotation, *i.e.*, by calculating the final NMI score. The simple nature of the modification of the label vector means that maximum values based on random splits can stay close to the scores obtained with the original assignment of four states (NMI score of 0.84, Figs. S3a and S5a–c).

The NetSAP analysis of the 4-state system with an average coherence of 0.7 showed high sensitivity, even to short 1-s input presentations (Fig. 6a). As a more realistic test, we thus wanted to explore the limits of NetSAP performance for a voltage recording with lower coherence and more states. Therefore, we carried out a 200-s simulation of a 50-neuron network with 20 states, a noise current of $I_n = 0.02 \mu\text{A}$ (mean coherence of 0.5), and an average residence time between 2000 ms and 3000 ms (Fig. 9). This test is particularly challenging as we did not impose any threshold on the minimum difference between inputs, and the neural states can thus overlap. The identification of states (as measured by an ARI of 0.54) is above chance levels (Fig. 9c–f). We wanted also to analyze if a higher number of neurons is beneficial for the correct identification of states. To do so, we performed 400 simulations of 100 s with a number of states ranging from 2 to 40, and

number of neurons ranging from 10 to 105. The parameters were kept identical to those in Fig. 5, with the exception of $I_n = 0.02 \mu\text{A}$ instead of $I_n = 0.01 \mu\text{A}$, and the average residence time which was randomly picked from a flat distribution between 1000 ms and 3000 ms. The resulting ARI scores are shown in Fig. 10a. As expected, a higher number of neurons permits a better identification of states. Conversely, NetSAP is not reliable when the number of states is larger than about 30. This failure is possibly due to overall simulation length, which does not allow for a sufficient level of recurrence when there are many states.

Next, we considered the case of multiple voltage patterns transmitted through different subpopulations of the network at the same time. For this purpose, we simulated a neural network consisting of two groups of neurons where each group represents a different voltage pattern. In practice, we made 10 simulations of 90 s each in which a subpopulation is involved in a 6-state system while the other is related to a different group of states (3-state system). To show that raw voltages are essential while firing rates would not be an effective pre-processing method, we simulated a network in which the two populations convey the voltage patterns synchronously, meaning that the theta oscillations are aligned. The other parameters were kept identical to Fig. 10a. We varied from 10 to 60 the number of neurons involved in the population representing the 6-state system and simulated a total of 70 neurons. The remaining neurons encoded the 3-state system. For the SAPHIRE analysis we imposed a true annotation with 3 or 6 states. The profile of the ARI value as a function of the subpopulation size increases for the networks with increasing amount of neurons that code for the system with the number of states used in the analysis (red dots in Fig. 10b). Qualitatively, an ARI score higher than 0.6 is obtained for the networks in which at least 40 of the 70 neurons are involved in the representation of the system with the same number of states used for SAPHIRE analysis.

Finally, we considered the possibility of having a subpopulation of completely unrelated neurons that spike at random times (Fig. 11). We modeled these neurons by removing the external oscillatory current I_w and increasing their noise current I_n by a factor of four. The other parameters were kept identical to the simulations used for Fig. 10a with 70 total neurons and 4 patterns, while we varied the number of unrelated neurons. The firing rates (Fig. 11c) differ from those of the original system (Fig. 2g) because the spiking times of the decoupled neurons are randomly distributed. The different voltage traces of the two populations are depicted in Fig. 11d. We simulated 22 neural networks with different numbers of unrelated neurons (from 2 to 65), and we calculated the resulting ARI (Fig. 11e). We found that the ARI score drops significantly when the number of random spiking neurons grows above 50% of the total of 70 neurons.

4. Discussion

We propose NetSAP as a data-driven analysis tool to identify neural states from voltage recordings. NetSAP was successfully applied to realistic albeit simplified neural simulations. With four different external inputs to the neural network, NetSAP was able to identify their four major corresponding neural states using only the raw voltage levels (Fig. 6a and b). Conversely, a standard clustering approach recognized only partially one of the states (k-means++ (Arthur and Vassilvitskii, 2007), Fig. 6b). A naive principal component analysis failed to isolate any of the four neural states (Fig. 7a).

In a more challenging test with 20 different patterns in 200 s of simulation (mean coherence of 0.5) NetSAP successfully mapped more than 65% of the snapshots to their correct labels, resulting in a total ARI of 0.54 (Fig. 9). The studied system was particularly challenging due to the low levels of recurrence and the short life times of states. In addition, we explored how the identification of states can be affected by the number of neurons and states involved (Fig. 10a). In the cases in which the number of states is high, a fixed number of snapshots limits the




Fig. 11. NetsAP analysis of simulations of a 70-neuron network in which a variable number of neurons spikes randomly and the remaining neurons receive input corresponding to four states. (a–c) A segment of 1 s is shown for a network that has half of its neurons spiking randomly. For illustrative purposes only 50 of the 70 neurons are shown. (b) To model completely unrelated neurons we modified the external input to the network I_{ext} by removing the oscillatory component and enhancing the colored noise (I_n is increased 4 fold). Here, the external current is shown for each neuron. (c) The firing rates are modified in comparison to Fig. 2g as the background neurons tend to disperse the theta oscillation due to their random spiking times. (d) Voltage levels of two theta oscillations after 5000 ms. In dark red we highlight the 25 neurons that spike randomly. (e) Using NetSAP, we analysed 22 simulations of 100 s each of a 70-neuron network with a variable number of random spiking neurons ranging from 2 to 65. All the other parameters are kept identical to Fig. 10a.

persistence time and recurrence of each state. As the SAPPHERE plot relies on the observed sampling density to detect states, it is likely that collecting and analyzing more snapshots would result in better scores. We also found that a higher number of information-carrying neurons allows NetSAP to achieve a better identification of states. To further evaluate the sensitivity of NetSAP on the number of neurons relevant to the state presentation we analyzed how the score drops if there is a mismatch between the true annotation and the predicted one (Fig. 10b) or if a subpopulation of neurons spikes randomly (Fig. 11e). The latter analysis provides evidence that a correct identification of states is possible if at least half of the neurons encode the signal of interest.

The main weakness in our simulation is its limited heterogeneity. An interesting improvement would be to enhance the number of neurons involved by closing the circuitry with a detailed excitatory population. Furthermore, it is possible to use a broader range of parameters to model each type of neuron. The enhancement of cellular heterogeneity can take advantage of known parameter clusters that are described experimentally (Tricoire et al., 2011).

Notably, the NetSAP work flow used here requires two essential hyperparameters: a similarity measure to infer the functional network and the distance metric between networks. For the latter we used the Euclidean distance without considering optimization. In NetSAP, we require another important parameter, *viz.*, the length of the time window, which was shown to be robust in the range 60–120 ms (Fig. S1d). It is also possible to overcome the limitation of a windowed approach, *i.e.*, the requirement to expect consistent time scales and artifacts due to noise factors, by transforming the data to, *e.g.*, the frequency domain. Among these methods, dynamic connectivity detection (Xu and Lindquist, 2015) and wavelet transform coherence (Chang and Glover, 2010), are cited here.

Even if we did not stress the selection of the number of states as a fundamental parameter, it is relevant for common experimental settings in which the ground truth (true annotation) is of course unknown. In these cases, the selected barriers could be inconsistent with available behavioural annotations. This is clearly possible because multi-neuronal data extracted from real networks tend to carry much more information than what is contained in a simple annotation. To implement some supervision in feature space, one could, for example, employ the NetSAP procedure to select neurons by a simple bootstrapping where the ARI scores highlight those neurons that are most consistent with a

given annotation. For this task, specific techniques like gradient optimization and simulated annealing are available as well (Satuvuori et al., 2018). This procedure would go beyond the usual experimental methods, which aim at average activity as their major selection rule (Ahrens et al., 2013; Cohen and Kohn, 2011; Harvey et al., 2012; Freeman et al., 2014; Lopes-dos Santos et al., 2011; Peyrache et al., 2009). The number of neurons, C , is also of practical importance for NetSAP. The initial inference of networks leads to N snapshots of dimensionality $D = C^2$, and these are fed into the dimensionality reduction by PCA. In a sparse regime, where N is not significantly larger than C , this may cause issues. However, as NetSAP relies on differences in sampling density for identifying states, this regime is poorly suited to begin with.

For real data sets, it is clearly of supreme interest to match a ground truth defined by a relevant behavioral annotation with neural states. As mentioned above, it is a frequent assumption that not all neurons in the network are relevant, and the identification of behaviorally critical ones is a common goal. For example, consider an experiment aimed at recognizing hallmarks of a specific neural disorder against a healthy control group. A direct analysis may rely on overall activity levels for individual neurons or on the correlation of this activity with the behavioral ground truth. Instead, the NetSAP pipeline offers an indirect route to prune the measured data sets: by bootstrapping in neuron space and a subsequent, unsupervised inference of neuronal states, it may be possible to find partitions that have significant NMI values with the behavioral annotation. This offers a natural route to not only prevent overfitting but also to independently estimate whether the behavioral state is clearly encoded in these neurons at all.

The NetSAP analysis pipeline can be improved further. For example, more sophisticated network inference methods could unveil complex relationships between neurons (*e.g.*, using transfer entropy (Schreiber, 2000) or directed information (Cai et al., 2017)). As alluded to above, the metric used for measuring similarity between inferred networks (snapshots) could be optimized using network-specific metrics like the Ipsen-Mikhailov one (Ipsen and Mikhailov, 2002).

Fundamental questions do remain, however. For example, is it possible that the distributed and transient nature of the neural code is simply too elusive for a detailed recognition of behavioural context? Even then, we believe that a combination of experiments aimed at discovering the real synaptic topology (structural connectivity) and a

behavioral NetSAP analysis (functional connectivity) could provide hints for how neurons exchange information. For these reasons, the NetSAP procedure has the potential to become a valuable tool in the analysis of multi-neuronal voltage recordings. Nevertheless, applications on experimental data will be essential in demonstrating NetSAP's usefulness to a wider community.

One area of particular interest to us is the analysis of neural systems affected by neurological disorders. In future work, we intend to explore how neural networks can (or cannot) act in compromised states. Consequently, applications will focus on highlighting aspects that play a key role in neural deregulation, e.g., in the case of epileptic seizure onset.

NetSAP is coded in CampaRi, an R package which is open source and already available on gitlab (<https://gitlab.com/CaflischLab/CampaRi>). This software is derived from the Fortran version of the progress index algorithm implemented in CAMPARI. (<http://campari.sourceforge.net/index.html>).

Conflict of interest statement

The authors declare no conflicts of interest.

Acknowledgements

We thank Francesco Cocina for many interesting discussions. This work was supported by grants of the Swiss National Science Foundation (to A. C.) and the Forschungskredit of the University of Zürich (to D. G.).

Appendix A. Supplementary data

Supplementary data associated with this article can be found, in the online version, at <https://doi.org/10.1016/j.jneumeth.2019.01.019>.

References

- Acsády, L., Arabadzisz, D., Freund, T.F., 1996a. Correlated morphological and neurochemical features identify different subsets of vasoactive intestinal polypeptide-immunoreactive interneurons in rat hippocampus. *Neuroscience* 73 (2), 299–315.
- Acsády, L., Görcs, T.J., Freund, T.F., 1996b. Different populations of vasoactive intestinal polypeptide-immunoreactive interneurons are specialized to control pyramidal cells or interneurons in the hippocampus. *Neuroscience* 73 (2), 317–334.
- Ahrens, M.B., Orger, M.B., Robson, D.N., Li, J.M., Keller, P.J., 2013. Whole-brain functional imaging at cellular resolution using light-sheet microscopy. *Nat. Methods* 10 (5), 413–420.
- Albanese, D., Filosi, M., Visintainer, R., Riccadonna, S., Jurman, G., Furlanello, C., 2013. minerva and minery: a C engine for the MINE suite and its r, python and MATLAB wrappers. *Bioinformatics* 29 (3), 407–408.
- Allen, E.A., Damaraju, E., Plis, S.M., Erhardt, E.B., Eichele, T., Calhoun, V.D., 2014. Tracking whole-brain connectivity dynamics in the resting state. *Cereb. Cortex* 24 (3), 663–676.
- Allen, E.A., Damaraju, E., Eichele, T., Wu, L., Calhoun, V.D., 2018. EEG signatures of dynamic functional network connectivity states. *Brain Topogr.* 31 (1), 101–116.
- Arthur, D., Vassilvitskij, S., 2007. k-means++: the advantages of careful seeding. *Proceedings of the Eighteenth Annual ACM-SIAM Symposium on Discrete Algorithms*, dl.acm.org 1027–1035.
- Atallah, B.V., Scanziani, M., 2009. Instantaneous modulation of gamma oscillation frequency by balancing excitation with inhibition. *Neuron* 62 (4), 566–577.
- Bacci, M., Vitalis, A., Caflisch, A., 2015. A molecular simulation protocol to avoid sampling redundancy and discover new states. *Biochim. Biophys. Acta* 1850 (5), 889–902.
- Baici, A., 2015. Kinetics of Enzyme-Modifier Interactions. Springer Vienna, Vienna.
- Belluscio, M.A., Mizuseki, K., Schmidt, R., Kempter, R., Buzsáki, G., 2012. Cross-frequency phase-phase coupling between θ and γ oscillations in the hippocampus. *J. Neurosci.* 32 (2), 423–435.
- Bezaire, M.J., Soltesz, I., 2013. Quantitative assessment of CA1 local circuits: knowledge base for interneuron-pyramidal cell connectivity. *Hippocampus* 23 (9), 751–785.
- Bialek, W., Rieke, F., de Ruyter van Steveninck, R.R., Warland, D., 1991. Reading a neural code. *Science* 252 (5014), 1854–1857.
- Blöchliger, N., Vitalis, A., Caflisch, A., 2013. A scalable algorithm to order and annotate continuous observations reveals the metastable states visited by dynamical systems. *Comput. Phys. Commun.* 184 (11), 2446–2453.
- Blöchliger, N., Vitalis, A., Caflisch, A., 2014. High-resolution visualisation of the states and pathways sampled in molecular dynamics simulations. *Sci. Rep.* 4, 6264.
- Blöchliger, N., Caflisch, A., Vitalis, A., 2015. Weighted distance functions improve analysis of high-dimensional data: application to molecular dynamics simulations. *J. Chem. Theory Comput.* 11 (11), 5481–5492.
- Blasco-Ibáñez, J.M., Freund, T.F., 1995. Synaptic input of horizontal interneurons in stratum oriens of the hippocampal CA1 subfield: structural basis of feed-back activation. *Eur. J. Neurosci.* 7 (10), 2170–2180.
- Boerlin, M., Denève, S., 2011. Spike-Based population coding and working memory. *PLoS Comput. Biol.* 7 (2), e1001080.
- Boerlin, M., Machens, C.K., Denève, S., 2013. Predictive coding of dynamical variables in balanced spiking networks. *PLoS Comput. Biol.* 9 (11), e1003258.
- Buckley, A.W., Scott, R., Tyler, A., Mahoney, J.M., Thurm, A., Farmer, C., Swedo, S., Burroughs, S.A., Holmes, G.L., 2015. State-dependent differences in functional connectivity in young children with autism spectrum disorder. *EBioMedicine* 2 (12), 1905–1915.
- Buzsáki, G., Leung, L.W., Vanderwolf, C.H., 1983. Cellular bases of hippocampal EEG in the behaving rat. *Brain Res.* 287 (2), 139–171.
- Buzsáki, G., 2001. Hippocampal GABAergic interneurons: a physiological perspective. *Neurochem. Res.* 26 (8–9), 899–905.
- Buzsáki, G., 2002. Theta oscillations in the hippocampus. *Neuron* 33 (3), 325–340.
- Buzsáki, G., 2004. Large-scale recording of neuronal ensembles. *Nat. Neurosci.* 7 (5), 446–451.
- Cai, Z., Neveu, C.L., Baxter, D.A., Byrne, J.H., Aazhang, B., 2017. Inferring neuronal network functional connectivity with directed information. *J. Neurophysiol.* 118 (2), 1055–1069.
- Campello, R.J.G.B., Moulavi, D., Sander, J., 2013. Density-based clustering based on hierarchical density estimates. *Advances in Knowledge Discovery and Data Mining*. Springer Berlin Heidelberg, pp. 160–172.
- Campello, R.J.G.B., Moulavi, D., Zimek, A., Sander, J., 2015. Hierarchical density estimates for data clustering, visualization, and outlier detection. *ACM Trans. Knowl. Discov. Data* 10 (1), 5.
- Chamberland, S., Topolnik, L., 2012. Inhibitory control of hippocampal inhibitory neurons. *Front. Neurosci.* 6, 165.
- Chang, C., Glover, G.H., 2010. Time-frequency dynamics of resting-state brain connectivity measured with fMRI. *Neuroimage* 50 (1), 81–98.
- Cocco, S., Monasson, R., Posani, L., Tavoni, G., 2017. Functional networks from inverse modeling of neural population activity. *Curr. Opin. Syst. Biol.* 3, 103–110.
- Cohen, M.R., Kohn, A., 2011. Measuring and interpreting neuronal correlations. *Nat. Neurosci.* 14 (7), 811–819.
- Cunningham, J.P., Yu, B.M., 2014. Dimensionality reduction for large-scale neural recordings. *Nat. Neurosci.* 17 (11), 1500–1509.
- Debanne, D., Bialowas, A., Rama, S., 2012. What are the mechanisms for analogue and digital signalling in the brain? *Nat. Rev. Neurosci.* 14, 63.
- Denève, S., Machens, C.K., 2016. Efficient codes and balanced networks. *Nat. Neurosci.* 19 (3), 375–382.
- Denk, W., Strickler, J.H., Webb, W.W., 1990. Two-photon laser scanning fluorescence microscopy. *Science* 248 (4951), 73–76.
- Du, J., Blanche, T.J., Harrison, R.R., Lester, H.A., Masmanidis, S.C., 2011. Multiplexed, high density electrophysiology with nanofabricated neural probes. *PLoS ONE* 6 (10), e26204.
- Dunn, T.W., Mu, Y., Narayan, S., Randlett, O., Naumann, E.A., Yang, C.-T., Schier, A.F., Freeman, J., Engert, F., Ahrens, M.B., 2016. Brain-wide mapping of neural activity controlling zebrafish exploratory locomotion. *Elife* 5, e12741.
- Engel, A.K., Fries, P., Singer, W., 2001. Dynamic predictions: oscillations and synchrony in top-down processing. *Nat. Rev. Neurosci.* 2 (10), 704–716.
- Ester, M., Kriegel, H.-P., Sander, J., Xu, X., et al., 1996. A density-based algorithm for discovering clusters in large spatial databases with noise. *Kdd*, vol. 96. pp. 226–231.
- Freeman, J., Vladimirov, N., Kawashima, T., Mu, Y., Sofroniew, N.J., Bennett, D.V., Rosen, J., Yang, C.-T., Looger, L.L., Ahrens, M.B., 2014. Mapping brain activity at scale with cluster computing. *Nat. Methods* 11 (9), 941–950.
- Freund, T.F., Antal, M., 1988. GABA-containing neurons in the septum control inhibitory interneurons in the hippocampus. *Nature* 336 (6195), 336170a0.
- Freund, T.F., Buzsáki, G., 1996. Interneurons of the hippocampus. *Hippocampus* 6 (4), 347–470.
- Freund, T.F., Gulyás, A.I., 1997. Inhibitory control of GABAergic interneurons in the hippocampus. *Can. J. Physiol. Pharmacol.* 75 (5), 479–487.
- Fries, P., 2015. Rhythms for cognition: communication through coherence. *Neuron* 88 (1), 220–235.
- Gallego, J.A., Perich, M.G., Miller, L.E., Solla, S.A., 2017. Neural manifolds for the control of movement. *Neuron* 94 (5), 978–984.
- Gentet, L.J., Avermann, M., Matyas, F., Staiger, J.F., Petersen, C.C.H., 2010. Membrane potential dynamics of GABAergic neurons in the barrel cortex of behaving mice. *Neuron* 65 (3), 422–435.
- Goodman, D., Brette, R., 2008. Brian: a simulator for spiking neural networks in python. *Front. Neuroinform.* 2, 5.
- Goodman, D.F.M., Stimberg, M., Yger, P., Brette, R., 2014. Brian 2: neural simulations on a variety of computational hardware. *BMC Neurosci.* 15 (1), 1.
- Grossberger, L., Battaglia, F.P., Vinck, M., 2018. Unsupervised clustering of temporal patterns in high-dimensional neuronal ensembles using a novel dissimilarity measure. *PLoS Comput. Biol.* 14 (7), e1006283.
- Gulyás, A.I., Görcs, T.J., Freund, T.F., 1990. Innervation of different peptide-containing neurons in the hippocampus by GABAergic septal afferents. *Neuroscience* 37 (1), 31–44.
- Gulyás, A.I., Hájós, N., Freund, T.F., 1996. Interneurons containing calretinin are specialized to control other interneurons in the rat hippocampus. *J. Neurosci.* 16 (10), 3397–3411.
- Harris, K.D., Csicsvari, J., Hirase, H., Dragoi, G., Buzsáki, G., 2003. Organization of cell assemblies in the hippocampus. *Nature* 424 (6948), 552–556.

- Harvey, C.D., Coen, P., Tank, D.W., 2012. Choice-specific sequences in parietal cortex during a virtual-navigation decision task. *Nature* 484 (7392), 62–68.
- Headley, D.B., Paré, D., 2017. Common oscillatory mechanisms across multiple memory systems. *npj Sci. Learn.* 2 (1), 1.
- Helmchen, F., Denk, W., 2005. Deep tissue two-photon microscopy. *Nat. Methods* 2 (12), 932–940.
- Henze, D.A., Borhegyi, Z., Csicsvari, J., Mamiya, A., Harris, K.D., Buzsáki, G., 2000. Intracellular features predicted by extracellular recordings in the hippocampus in vivo. *J. Neurophysiol.* 84 (1), 390–400.
- Hindriks, R., Adhikari, M.H., Murayama, Y., Ganzetti, M., Mantini, D., Logothetis, N.K., Deco, G., 2016. Can sliding-window correlations reveal dynamic functional connectivity in resting-state fMRI? *Neuroimage* 127, 242–256.
- Hirayama, J.-I., Ogawa, T., Hyvärinen, A., 2015. Unifying blind separation and clustering for resting-state EEG/MEG functional connectivity analysis. *Neural Comput.* 27 (7), 1373–1404.
- Hodgkin, A.L., Huxley, A.F., 1952. A quantitative description of membrane current and its application to conduction and excitation in nerve. *J. Physiol.* 117 (4), 500–544.
- Hubert, L., Arabie, P., 1985. Comparing partitions. *J. Classif.* 2 (1), 193–218.
- Huettel, S.A., Song, A.W., McCarthy, G., et al., 2004. Functional Magnetic Resonance Imaging, vol. 1 Sinauer Associates Sunderland, MA.
- Hutchison, R.M., Womelsdorf, T., Allen, E.A., Bandettini, P.A., Calhoun, V.D., Corbetta, M., Della Penna, S., Duyn, J.H., Glover, G.H., Gonzalez-Castillo, J., Handwerker, D.A., Keilholz, S., Kiviniemi, V., Leopold, D.A., de Pasquale, F., Sporns, O., Walter, M., Chang, C., 2013. Dynamic functional connectivity: promise, issues, and interpretations. *Neuroimage* 80, 360–378.
- Ipsen, M., Mikhailov, A.S., 2002. Evolutionary reconstruction of networks. *Phys. Rev. E Stat. Nonlinear Soft Matter Phys.* 66 (4 Pt 2), 046109.
- Kass, R.E., Eden, U.T., Brown, E.N., 2014. Analysis of Neural Data. Springer.
- Kirino, E., Tanaka, S., Fukuta, M., Inami, R., Arai, H., Aoki, S., 2017. Simultaneous resting-state functional MRI and electroencephalography recordings of functional connectivity in patients with schizophrenia. *Psychiatry Clin. Neurosci.* 71 (4), 262–270.
- Kirst, C., Timme, M., Battaglia, D., 2016. Dynamic information routing in complex networks. *Nat. Commun.* 7, 11061.
- Kollmorgen, S., Hahnloser, R.H.R., 2014. Dynamic alignment models for neural coding. *PLoS Comput. Biol.* 10 (3), e1003508.
- Krivov, S.V., Karplus, M., 2006. One-dimensional free-energy profiles of complex systems: progress variables that preserve the barriers. *J. Phys. Chem. B* 110 (25), 12689–12698.
- Lewis, D.A., Hashimoto, T., Volk, D.W., 2005. Cortical inhibitory neurons and schizophrenia. *Nat. Rev. Neurosci.* 6 (4), 312–324.
- Li, N., Daie, K., Svoboda, K., Druckmann, S., 2016. Robust neuronal dynamics in premotor cortex during motor planning. *Nature* 532 (7600), 459–464.
- Liao, W.T., 2005. Clustering of time series data—a survey. *Pattern Recognit.* 38 (11), 1857–1874.
- Lim, S., Goldman, M.S., 2013. Balanced cortical microcircuitry for maintaining information in working memory. *Nat. Neurosci.* 16, 1306.
- Lim, S., Goldman, M.S., 2014. Balanced cortical microcircuitry for spatial working memory based on corrective feedback control. *J. Neurosci.* 34 (20), 6790–6806.
- Lisman, J.E., Coyle, J.T., Green, R.W., Javitt, D.C., Benes, F.M., Heckers, S., Grace, A.A., 2008. Circuit-based framework for understanding neurotransmitter and risk gene interactions in schizophrenia. *Trends Neurosci.* 31 (5), 234–242.
- Lopes-dos Santos, V., Conde-Ocazonez, S., Nicolelis, M.A.L., Ribeiro, S.T., Tort, A.B.L., 2011. Neuronal assembly detection and cell membership specification by principal component analysis. *PLoS ONE* 6 (6), e20996.
- Marín, O., 2012. Interneuron dysfunction in psychiatric disorders. *Nat. Rev. Neurosci.* 13 (2), 107–120.
- Marx, M., Haas, C.A., Häusser, U., 2013. Differential vulnerability of interneurons in the epileptic hippocampus. *Front. Cell. Neurosci.* 7, 167.
- McCormick, D.A., Connors, B.W., Lighthall, J.W., Prince, D.A., 1985. Comparative electrophysiology of pyramidal and sparsely spiny stellate neurons of the neocortex. *J. Neurophysiol.* 54 (4), 782–806.
- McInnes, L., Healy, J., Astels, S., 2017. HdbSCAN: hierarchical density based clustering. *J. Open Source Softw.* 2 (11), 205.
- McKay, B.M., Oh, M.M., Disterhoft, J.F., 2013. Learning increases intrinsic excitability of hippocampal interneurons. *J. Neurosci.* 33 (13), 5499–5506.
- Najafi, K., Ji, J., Wise, K.D., 1990. Scaling limitations of silicon multichannel recording probes. *IEEE Trans. Biomed. Eng.* 37 (1), 1–11.
- Ota, K., Omori, T., Aonishi, T., 2009. MAP estimation algorithm for phase response curves based on analysis of the observation process. *J. Comput. Neurosci.* 26 (2), 185–202.
- Pakhira, M.K., Bandyopadhyay, S., Maulik, U., 2004. Validity index for crisp and fuzzy clusters. *Pattern Recognit.* 37 (3), 487–501.
- Palmigiano, A., Geisel, T., Wolf, F., Battaglia, D., 2017. Flexible information routing by transient synchrony. *Nat. Neurosci.* 20 (7), 1014–1022.
- Paninski, L., 2004. Maximum likelihood estimation of cascade point-process neural encoding models. *Network* 15 (4), 243–262.
- Pelkey, K.A., Chittajallu, R., Craig, M.T., Tricoire, L., Wester, J.C., McBain, C.J., 2017. Hippocampal GABAergic inhibitory interneurons. *Physiol. Rev.* 97 (4), 1619–1747.
- Peyrache, A., Khamassi, M., Benchenane, K., Wiener, S.I., Battaglia, F.P., 2009. Replay of rule-learning related neural patterns in the prefrontal cortex during sleep. *Nat. Neurosci.* 12 (7), 919–926.
- Pitkow, X., Angelaki, D.E., 2017. Inference in the brain: statistics flowing in redundant population codes. *Neuron* 94 (5), 943–953.
- Preti, M.G., Bolton, T.A., Van De Ville, D., 2017. The dynamic functional connectome: state-of-the-art and perspectives. *Neuroimage* 160, 41–54.
- Reshef, D.N., Reshef, Y.A., Finucane, H.K., Grossman, S.R., McVean, G., Turnbaugh, P.J., Lander, E.S., Mitzenmacher, M., Sabeti, P.C., 2011. Detecting novel associations in large data sets. *Science* 334 (6062), 1518–1524.
- Rolls, E.T., Treves, A., 2011. The neuronal encoding of information in the brain. *Prog. Neurobiol.* 95 (3), 448–490.
- Romo, R., de Lafuente, V., 2013. Conversion of sensory signals into perceptual decisions. *Prog. Neurobiol.* 103, 41–75.
- Satuvuori, E., Mulansky, M., Daffertshofer, A., Kreuz, T., 2018. Using spike train distances to identify the most discriminative neuronal subpopulation. *J. Neurosci. Methods* 308, 354–365.
- Schaub, M.T., Schultz, S.R., 2012. The Ising decoder: reading out the activity of large neural ensembles. *J. Comput. Neurosci.* 32 (1), 101–118.
- Schneidman, E., Berry 2nd, M.J., Segev, R., Bialek, W., 2006. Weak pairwise correlations imply strongly correlated network states in a neural population. *Nature* 440 (7087), 1007–1012.
- Schreiber, T., 2000. Measuring information transfer. *Phys. Rev. Lett.* 85 (2), 461–464.
- Sengupta, B., Laughlin, S.B., Niven, J.E., 2014. Consequences of converting graded to action potentials upon neural information coding and energy efficiency. *PLoS Comput. Biol.* 10 (1), e1003439.
- Shadlen, M.N., Newsome, W.T., 1998. The variable discharge of cortical neurons: implications for connectivity, computation, and information coding. *J. Neurosci.* 18 (10), 3870–3896.
- Stefanelli, T., Bertolini, C., Lüscher, C., Müller, D., Mendez, P., 2016. Hippocampal somatostatin interneurons control the size of neuronal memory ensembles. *Neuron* 89 (5), 1074–1085.
- Stevenson, I.H., Kording, K.P., 2011. How advances in neural recording affect data analysis. *Nat. Neurosci.* 14 (2), 139–142.
- Stosiek, C., Garaschuk, O., Holthoff, K., Konnerth, A., 2003. In vivo two-photon calcium imaging of neuronal networks. *Proc. Natl. Acad. Sci. USA* 100 (12), 7319–7324.
- Strehl, A., Ghosh, J., 2002. Cluster ensembles – a knowledge reuse framework for combining multiple partitions. *J. Mach. Learn. Res.* 3 (December), 583–617.
- Tóth, K., Freund, T.F., Miles, R., 1997. Disinhibition of rat hippocampal pyramidal cells by GABAergic afferents from the septum. *J. Physiol.* 500 (Pt 2), 463–474.
- Trevelyan, A.J., Sussillo, D., Yuste, R., 2007. Feedforward inhibition contributes to the control of epileptiform propagation speed. *J. Neurosci.* 27 (13), 3383–3387.
- Tricoire, L., Pelkey, K.A., Erkkila, B.E., Jeffries, B.W., Yuan, X., McBain, C.J., 2011. A blueprint for the spatiotemporal origins of mouse hippocampal interneuron diversity. *J. Neurosci.* 31 (30), 10948–10970.
- Truccolo, W., Eden, U.T., Fellous, M.R., Donoghue, J.P., Brown, E.N., 2005. A point process framework for relating neural spiking activity to spiking history, neural ensemble, and extrinsic covariate effects. *J. Neurophysiol.* 93 (2), 1074–1089.
- Tyan, L., Chamberland, S., Magnin, E., Camiré, O., Francavilla, R., David, L.S., Deisseroth, K., Topolnik, L., 2014. Dendritic inhibition provided by interneuron-specific cells controls the firing rate and timing of the hippocampal feedback inhibitory circuitry. *J. Neurosci.* 34 (13), 4534–4547.
- van Vreeswijk, C., Sompolinsky, H., 1996. Chaos in neuronal networks with balanced excitatory and inhibitory activity. *Science* 274 (5293), 1724–1726.
- L. Vendramin, R. J. G. B. Campello, E. R. Hruschka, Relative clustering validity criteria: a comparative overview, Stat. Anal. Data Min. 31.**
- Vidaurre, D., Quinn, A.J., Baker, A.P., Dupret, D., Tejero-Cantero, A., Woolrich, M.W., 2016. Spectrally resolved fast transient brain states in electrophysiological data. *Neuroimage* 126, 81–95.
- Vinh, N.X., Epps, J., Bailey, J., 2010. Information theoretic measures for clusterings comparison: variants, properties, normalization and correction for chance. *J. Mach. Learn. Res.* 11 (October), 2837–2854.
- Vitalis, A., Cafilisch, A., 2012. Efficient construction of mesostate networks from molecular dynamics trajectories. *J. Chem. Theory Comput.* 8 (3), 1108–1120.
- Wang, X.J., Buzsáki, G., 1996. Gamma oscillation by synaptic inhibition in a hippocampal interneuronal network model. *J. Neurosci.* 16 (20), 6402–6413.
- Wang, J.W., Wong, A.M., Flores, J., Voshall, L.B., Axel, R., 2003. Two-photon calcium imaging reveals an odor-evoked map of activity in the fly brain. *Cell* 112 (2), 271–282.
- Welch, P., 1967. The use of fast Fourier transform for the estimation of power spectra: a method based on time averaging over short, modified periodograms. *IEEE Trans. Audio Electroacoust.* 15 (2), 70–73.
- Xu, Y., Lindquist, M.A., 2015. Dynamic connectivity detection: an algorithm for determining functional connectivity change points in fMRI data. *Front. Neurosci.* 9, 285.
- Xue, M., Atallah, B.V., Scanziani, M., 2014. Equalizing excitation-inhibition ratios across visual cortical neurons. *Nature* 511 (7511), 596–600.
- Yaesoubi, M., Allen, E.A., Miller, R.L., Calhoun, V.D., 2015a. Dynamic coherence analysis of resting fMRI data to jointly capture state-based phase, frequency, and time-domain information. *Neuroimage* 120, 133–142.
- Yaesoubi, M., Miller, R.L., Calhoun, V.D., 2015b. Mutually temporally independent connectivity patterns: a new framework to study the dynamics of brain connectivity at rest with application to explain group difference based on gender. *Neuroimage* 107, 85–94.
- Yu, J., Ferster, D., 2010. Membrane potential synchrony in primary visual cortex during sensory stimulation. *Neuron* 68 (6), 1187–1201.
- Yuste, R., 2015. From the neuron doctrine to neural networks. *Nat. Rev. Neurosci.* 16 (8), 487–497.