Statistical analysis of synchrony and synchrony propagation in massively parallel spike trains

Von der Fakultät für Mathematik, Informatik und Naturwissenschaften
der RWTH Aachen University zur Erlangung des akademischen Grades
eines Doktors der Naturwissenschaften genehmigte Dissertation

vorgelegt von

Emiliano Torre,
M.Sc. in Mathematics,
aus Torino, Italien

Berichter: Universitätsprofessorin Dr. rer. nat. Sonja Grün
Universitätsprofessor Dr. rer. nat. Bjorn Kampa
Professorin Dr. rer. nat. Laura Sacerdote

Tag der mündliche Prüfung: April 29, 2016

To my mother, Silvana

Abstract

The mechanisms by which the brain represents, processes and transmits information in terms of spiking activity are largely unknown. The coordination of spikes among neurons at millisecond temporal scale has been hypothesized as a mechanism of information coding. Spike synchronization is considered a signature of such cooperative dynamics. Experimental studies have found evidence of the occurrence of spike synchronization beyond chance level in relation to behavior. However, these findings were mostly limited to data from a small number of simultaneously recorded neurons, and were concluded on the basis of pairwise or low-order correlation analyses unable to conclude on the presence of neuronal interactions of any order.

Recent advances in electrophysiological technology enable recordings of hundreds of neurons in parallel, a number which is destined to grow. The increase in data dimensionality makes traditional methods for spike correlation analysis inapplicable due to the combinatorical explosion in the number of patterns, which raises severe computational and statistical problems. New tools for the analysis of big data are thus needed.

This essay delivers two novel methods for the analysis of spike synchronization in massively parallel spike trains. The first method, called SPADE (spike pattern detection and evaluation) is designed to perform an optimal search for synchronous spike patterns using efficient data mining techniques. The detected patterns are then tested for significance using a hierarchy of statistical tests that circumvents multiple testing issues.

We employ SPADE to investigate the occurrence of spike synchronization in the motor cortex of two macaque monkeys performing an instructed-delay reach-to-grasp task. The task consists of different epochs and was performed in different ways. By analyzing each epoch and trial type separately, we are able to associate the emergence of individual patterns of synchronous spikes to specific behaviors, and we investigate the spatial organization of synchronous spiking neurons in motor cortex. Our analysis reveals groups of neurons, aligned along a preferential medio-lateral orientation, which selectively synchronize their activity in relation to behavior.

We finally introduce a second statistical method, called ASSET (analysis of sequences of synchronous events), that searches for temporal sequences of synchronous spike patterns. Synchrony propagation has been long suggested as a plausible mechanism of information coding and transmission, successfully implemented for instance in the synfire chain model in a manner compatible with electrophysiological constraints. Here we provide, to the best of our knowledge, the first statistical tool to detect sequential synchronous events in massively parallel spike trains.

In summary, this work delivers two novel powerful statistical tools, SPADE and ASSET, to analyze spike synchrony and synchrony propagation in parallel spike trains. These methods are designed for application to large data from hundreds or thousands of simultaneous recordings, and cope with typical features of real data, such as non-stationarity of firing rates in time and across neurons or regularity in the spiking activity. The application of SPADE to motor cortex data from two macaque monkeys performing a reach-to-grasp task allows us to shed new light on the spatial and temporal organization of spike synchronization in the cerebral cortex in relation to behavior.

Zusammenfassung

Die Mechanismen des Gehirns zur Darstellung, Verarbeitung und Übertragung von Information sind weitgehend unbekannt. Ein möglicher Weg zur Informationscodierung ist die koordinierte Aussendung von Impulsen ("Spikes") zwischen Neuronen mit Präzision im Millisekundenbereich. Die Synchronisation von Spikes wird als Anzeichen solch kooperativer Dynamik angesehen. Experimentelle Studien haben Belege von Spikesynchronisation über Zufallsniveau gefunden, jedoch beruhen die meisten dieser Entdeckungen auf kleinen Datenmengen mit wenigen gleichzeitig beobachteten Neuronen. Außerdem basieren diese Entdeckungen auf der Analyse paarweiser oder niedrigrangiger Korrelationen, die das Auftreten von neuronalen Wechselwirkungen höherer Ordnung vernachlässigen.

Fortschritte in der elektrophysiologischen Technologie bieten seit kurzem die Möglichkeit die Aktivität hunderter Neurone gleichzeitig aufzuzeichnen, eine Zahl, die noch weiter ansteigen wird. Die anwachsende Dimensionalität der Daten macht die traditionellen Methoden für Spikekorrelationsanalysen aufgrund der kombinatorischen Explosion der Zahl von Kombinationen, die statistische Probleme aufwirft und die Rechnungen erschwert, ungeeignet. Aufgrund dessen werden neue Werkzeuge zur Analyse großer Datensätze benötigt.

Diese Arbeit liefert zwei neue Verfahren zur Analyse der Synchronisation in massiv parallelen Spikesequenzen. Die erste Methode, namens SPADE (Spike Pattern Detection and Evaluation), wurde entwickelt um eine optimale Suche nach Mustern von synchronen Spikes durch effiziente Data-Miningtechniken durchzuführen. Die auf diese Weise gefundenen Muster werden dann auf ihre Signifikanz unter Verwendung einer Hierarchie von statistischen Tests, die das Problem der Alphafehler-Kumulierung ("multiple Testing") vermeiden, bewertet.

Wir verwenden SPADE um das Auftreten von Spikesynchronisation in Daten aus dem Motorcortex von zwei Affen zu untersuchen, während die Tieren eine Aufgabe ausführen, bei der sie einen Gegenstand erreichen und greifen. Die Aufgabe besteht aus verschiedenen Phasen (Epochen) und wurde in verschiedenen Varianten wiederholt. Durch die separate Analyse von verschiedenen Epochen und Versuchsvarianten, verbinden wir jedes Muster synchroner Spikes mit einem bestimmten Verhalten und untersuchen die räumliche Organisation synchron aktiver Neuronen im Motorcortex.

Schließlich führen wir eine zweite Methode namens ASSET (Analysis of Sequences of Synchronous Events) ein, welche zeitliche Sequenzen synchroner Spikemuster detektiert. Die Ausbreitung von synchroner Aktivität wird seit langem als ein plausibler Mechanismus der Informationscodierung und übertragung betrachtet, erfolgreich implementiert zum Beispiel im "synfire chain" Modell, kompatibel mit den elektrophysiologischen Rahmenbedingungen. Wir präsentieren ein neues statistisches Werkzeug, um sequentielle synchrone Ereignisse in Daten aus massiv parallelen Spikesequenzen zu entdecken.

Diese Methoden sind für die Anwendung auf Daten von hunderten oder tausenden gleichzeitig gemessenen Neuronen konzipiert und mit typischen Eigenschaften von realen Messungen wie Nicht-Stationarität von Feuerraten und Regelmäßigkeiten von Spikeaktivität. Die Anwendung von SPADE auf Daten vom Motorkortex zweier Affen gibt Aufschluss über auf die räumliche und zeitliche Organisation von Spikesynchronisation in der Großhirnrinde im Bezug auf Verhalten.

Contents

Chapter 1

Introduction

This essay presents two novel statistical methods to reveal patterns of synchronous spikes and temporal sequences of such patterns in high-dimensional data composed of massively parallel spike trains (MPST). The two methods are validated through application to test data covering a range of scenarios typically encountered in the analysis of electrophysiological data. We also show an application to experimental data from behaving monkeys, which allows us to analyze the emergence and organization of synchronous spiking in relation to behavior. In this chapter we discuss the relevance of spike synchronization in the neuronal circuitry, we provide a brief overview of existing methods for the analysis of synchrony in parallel spike trains, and we highlight their limitations in applications to high-dimensional data sets. These limitations motivate the need for new techniques that can be applied to MPST without incurring into unsoluble computational and statistical problems.

1.1 Neurons and spikes

The brain is the organ of the nervous system responsible for information processing in most animals. It receives information from the rest of the body, such as sensory stimuli, elaborates this information and guides responses by the rest of the body. The brain of primates comprises billions of cells of diverse types, wired together by thousands of billions of synapses in a complex network structure (Braitenberg and Schüz, 1998). Two broad classes of cells in the brain are neurons and glia cells (Kandel et al., 1991). Glia cells accomplish a number of tasks that support neurons and the brain tissue in their functioning, providing for instance nutrition and structural support, serving in the removal of dead cells and the reparation of the brain and the spinal cord, regulating the production or neurotransmitters. Neurons are excitable cells that receive and transmit electrical impulses known as action potentials, or *spikes*. Spikes generated in one neuron travel down its axon to the synapses connecting the cell to a post-synaptic neuron. Input spikes to the post-synaptic neuron generate post-synaptic currents (PSCs) that summate to trigger a new spike.

Figure 1.1 illustrates the time course of a typical spike. A spike consists of a rapid excursion of the membrane potential, i.e. of the voltage gradient between the outside and the inside of the cell, from the cell's "resting" state. In absence of electrical currents, the external medium in which a neuron is immersed contains more positive-charged ions than the cell's inside, determining a negative difference between the inside and the outside of the cell. The voltage difference is maintained to a *resting potential* by the cell's membrane, a barrier which has in resting condition low permeability to ions. Input currents to a neu-

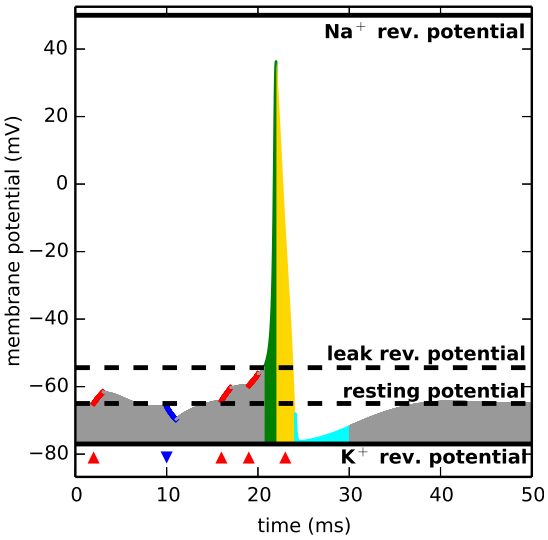

Figure 1.1: **Spike generation from summation of PSCs.** Input currents to a neuron travel down the neuron's dendrite to the soma. Excitatory/inhibitory inputs (red/blue triangles) depolarize/hyperpolarize the membrane potential (filled area). Depolarization to leak reversal potential triggers a spike, i.e. a rapid excursion to positive values (green area) followed by hyperpolarization of similar speed (yellow area, hard refractory period) during which excitatory input has no depolarization effect. Hyperpolarization typically undershoots the resting potential, causing a soft refractory period (cyan) followed by stabilization to resting potential.

rons can change this situation, by causing the opening of *ion channels* which let specific ions pass through, and thus by depolarizing the membrane potential towards less negative values or conversely by hyper-polarizing it. The most typical spike-generation mechanism consists in the opening of voltage-gated Na^+ channels as a consequence of sufficient depolarization determined by incoming input. Because Na^+ has higher concentration outside than inside the cell, when these channels open Na^+ ions start flowing inside the cell and depolarize the potential even faster, until the gradient reverses to even positive values. At this stage Na^+ channels inactivate due to sensitivity to the voltage, the situation reverses and the membrane potential hyperpolarizes again. The concomitant opening of voltage-gated K^+ channels allows an outflow of K^+ ions, further contributing to hyper-polarizing the membrane potential to values even below the resting potential. After these channels close, the potential stabilizes to reversal potential. A spike thus consists of three main successive phases, illustrated in Figure 1.1: i) Na^+-mediated fast ($\approx 1\,\text{ms}$) depolarization (green area), ii) fast ($1-2\,\text{ms}$) hyper-polarization, during which Na^+ channel are inactive and generation of new spikes is not possible (hard refractory period, yellow area), and iii) soft refractory period, lasting up to several ms, during which the activation of K^+ channels has an hyper-polarization effect on the membrane and the generation of new spikes is less likely than in resting condition (cyan).

1.2 Rate coding versus temporal coding

Spikes form the alphabet used by neurons to communicate among each other and with nerves and muscles. The discovery of spikes (Adrian, 1926b) and of the first mechanisms responsible for their generation and transmission (Hodgkin and Huxley, 1952b,a), led to the foundation of the modern understanding of spike-based information processing in the brain. Nevertheless, how spikes, as letters of this alphabet, are combined together to form the vocabulary used by neurons to communicate is still nowadays to a large extent unknown. Various spike-based paradigms of information coding have been proposed on the basis of experimental studies. Among them, two prominent but contrasting views are the rate coding and the temporal coding paradigms.

Rate coding assumes that information is represented and processed in terms of the modulation and co-modulation of the rate by which neurons emit spikes over time (*firing*

rate). Firing rate modulation has been one of the earliest observations following the discovery of spikes (Adrian and Zotterman, 1926; Adrian, 1926a), found almost ubiquitously across different systems (see e.g. Georgopoulos et al., 1986; Lamme and Spekreijse, 1998; Roelfsema et al., 2004). Rate-coding represents the input to a neuron in terms of rate of incoming spikes, computed by integrating over a temporal window of tens or hundreds of milliseconds. This typically results in a multitude of spikes, which makes the computation weakly sensitive to the presence or the precise occurrence time of single spikes and therefore robust to input noise (Shadlen and Newsome, 1994, 1995; London et al., 2010). However, a number of studies also suggests that in several circumstances this scheme would be too slow to process and transmit information (Gautrais and Thorpe, 1998; VanRullen et al., 2005) as well as being unable to provide unique representations of stimuli (Gerstner et al., 1997).

In contrast, temporal coding assumes than neurons coordinate the timing of their spikes on a millisecond time scale, for instance by synchronizing their activity or by firing in specific temporal sequences (Abeles, 1991). Hebb (1949) proposed that neurons in the brain represent and process information by organizing in functional groups, termed cell assemblies, rather than acting as individual entities. According to this paradigm, individal spikes and their precise timing rather than their counts over large temporal windows matter (Abeles, 1982b; König et al., 1996; Singer et al., 1997; Harris, 2005), making this scheme faster and more efficient than rate-coding (Gerstner et al., 1997; VanRullen et al., 2005; Gautrais and Thorpe, 1998) but also more sensitive to small perturbations in the input and therefore less robust to noise (see e.g. London et al., 2010). Spike synchrony is understood as a conceptually different mechanism of information processing than rate (co-)modulation, because it corresponds to reliable (probability 1) rather than stochastic (probability < 1) neuron activation (De la Rocha et al., 2007; Shea-Brown et al., 2008; Staude et al., 2008; Schultze-Kraft et al., 2013).

Millisecond precision in the arrival time of input spikes plays a critical role with respect to at least two features of the mechanism of spike generation. The first feature is the leaky nature of neuronal membranes. PSCs caused by pre-synaptic spikes decay over time, eventually vanishing. Thus, synchronous input spikes are more effective in triggering a spike response than asynchronous inputs, because the PSCs they elicit summate to threshold before decaying (Kandel et al., 1991, chapter 12). The second feature relates to a fundamental mechanism of synaptic adaptation known as spike-time dependent plasticity (STDP). STDP consists in the long-lasting potentiation (Bliss and Lomo, 1973; Lømo, 2003) or depotentiation (Stent, 1973; Massey and Bashir, 2007) of synaptic efficacy (i.e. the increase or decrease in the amplitude of PSCs caused by pre-synaptic spikes) depending on the relative timing of the output response with respect to the input (Stuart and Sakmann, 1994; Markram et al., 1997; Zhang et al., 1998). Thus, a post-synaptic spike following/preceeding a pre-synaptic spike by a few milliseconds can cause a substantial short-term strengthening/weakening in the amplitude of PSCs elicited by pre-synaptic spikes. The sign and amplitude of this change is highly sensitive to small changes in the delay between the pre- and the post-synaptic spike times (for a review, see Markram et al., 2011), as illustrated in Figure 1.2. The discovery of STDP mechanisms confirmed the decade-old prediction by Hebb (1949) that precisely timed input-output delays shape the synaptic connections in the neuronal circuitry by strengthening or weakening the post-synaptic response to pre-synaptic spikes.

Neurons with leaky membranes and whose synapses undergo STDP mechanisms may behave like coincidence detectors (Abeles, 1982b), able to sense and respond to synchronous input more reliably than to asynchronous input. Experimental studies provide evidence of the existence of coincidence detectors (Roy and Alloway, 2001; Bender et al., 2006; Fino

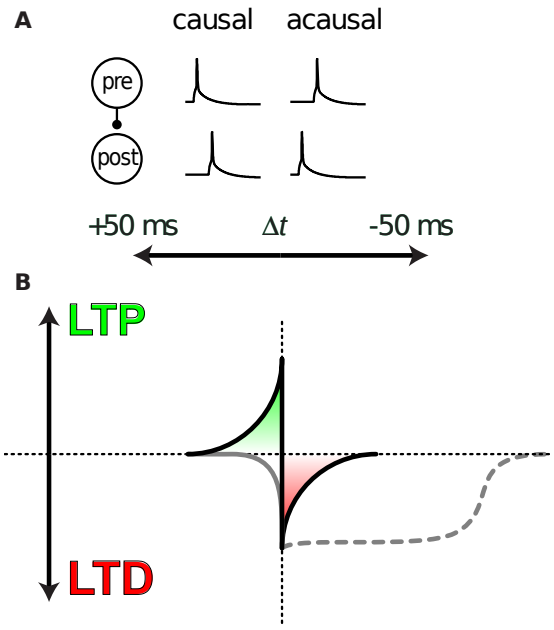

Figure 1.2: **Illustration of spike timing dependent plasticity.** The relative spike timing of the pre- and post-synaptic neurons determines changes in their synaptic efficacy. A pre-synaptic neuron repeatedly firing some milliseconds before/after the post-synaptic cell (causal/acausal activity; panel A) causes a long term potentiation/depotentiation of synapstic efficacy (panel B). At some cortical synapses, the temporal window for LTD is extended (dashed gray line; see Feldman, 2000; Sjöström et al., 2001). Figure reproduced from Figure 1 of Markram et al. (2011).

et al., 2010) and relate them to different encoding and decoding schemes Perez-Orive et al. (2004); Hong et al. (2012).

Spike synchronization has been shown to occur in relation to stimuli or behavior in various systems, such as the auditory cortex (Seki and Eggermont, 2003; Carr, 2004; Eggermont, 2015), the retina (Van Rullen and Thorpe, 2001; Hu and Bloomfield, 2003; Shlens et al., 2006; Pillow et al., 2008), the visual cortex (von der Malsburg, 1986; Engel et al., 1992; van der Togt et al., 2006; Berger et al., 2007; Smith and Kohn, 2008; Martin and von der Heydt, 2015), the motor cortex (Riehle et al., 1997; Baker et al., 2001; Shimazaki et al., 2012), the olivocerebellar tract (Lang, 2002; De Gruijl et al., 2014), the somatosensory cortex (Steinmetz et al., 2000; Reed et al., 2008; Harvey et al., 2013), the hyppocampus (Sakurai, 1996; Diba et al., 2014), the frontal and pre-frontal cortex (Vaadia et al., 1995; Sakurai and Takahashi, 2006; Fujisawa et al., 2008; Pipa and Munk, 2011), involving excitatory as well as inhibitory (Diba et al., 2014; Hu and Agmon, 2015) neurons. The diversity of species (in the cited references: cats, rats, monkeys, birds), systems and neuron types exhibiting behavior-related spike synchronization shows that rate coding may indeed be a valid mechanism of information processing in the neuronal circuitry.

Rate coding and temporal coding are two somewhat opposing paradigms of information processing, each with its own advantages and disadvantages. Convincing arguments and clear experimental evidence have been provided for both schemes. Nowadays, it seems plausible that both codes may be used by the brain (Kumar et al., 2010; Ainsworth et al., 2012), depending for instance on the cognitive process that takes place (Ahissar et al., 2000; Prescott and Sejnowski, 2008), the spatial distance between the neurons involved (Smith and Kohn, 2008), or the characteristics of the synaptic input (Rudolph and Destexhe, 2003).

1.3 Existing methods for the analysis of spike synchrony and their shortcomings

A variety of methods exist for the analysis of synchrony in parallel spike trains. The majority of these methods use a statistical approach, by testing the null hypothesis that neurons are mutually independent. If the null hypothesis is rejected, for instance because

the observed synchrony exceeds the level expected under independence, one accepts the alternative hypothesis that correlations exist and that the population under observation indeed synchronizes its activity. In order to disentangle rate-coding and temporal-coding dynamics and correctly determine the level of synchronization expected under the null hypothesis assuming a pure rate-coding scheme, firing rates and their modulation over time have to be incorporated in the formulation of the null distribution. Here we review existing statistical methods for the analysis of synchrony in parallel spike trains.

Most of these methods have been introduced to discover correlations of low order (i.e. between two or few neurons) and/or in data of limited size, consisting of a small number of neurons recorded simultaneously. Recent advances in electrophysiological technology enable hundreds of simultaneous recordings (Buzsaki, 2004), a number destined to grow fast to thousands or beyond in the next years (Schwarz et al., 2014). We show that classical methods cannot be applied to data of this size in a straightforward way, due to severe computational or statistical problems that the higher dimensionality raises. This limitation motivates the need for new techniques applicable to next-generation recordings, as the two methods we propose in the next chapters.

1.3.1 Cross-correlation histogram

The first technique to measure the correlation between two spike trains was proposed by Perkel et al. (1967b), who introduced the *cross-correlation histogram* (CCH) to measure the tendency of a target neuron to emit a spike a certain time before or after another reference neuron. The CCH counts, for any delay interval $\Theta = [\tau, \tau + \delta]$, the number of spikes of the target neuron preceeding or following any spike of the target neuron by an amount $\theta \in \Theta$ (see Figure 1.3a).

For infinitely (in practice: sufficiently) long spike trains with Poisson statistics and stationary firing rates, spike train independence yields flat a CCH. If the target neuron tends to fire a preferred time τ before or after the reference neuron, a peak appears in their CCH at that delay. Synchronization between the two neurons yields a peak around delay 0. Peaks in the CCH however may be caused not only by (delayed) synchrony, but also by other features of the activity of the individual neurons, such as high inter-spike interval (ISI) regularity - which generates oscillatory CCHs at a frequency inversely proportional to the preferred ISI length - or modulation of the neuronal firing rates. Various analytical and numerical methods have been introduced to test the statistical significance of CCH peaks given these and other features of the data (for a review, see Grün and Rotter, 2010, chapters 6 and 17), as illustrated in Figure 1.3b.

The CCH is not designed to target synchrony only, but can reveal delayed correlations as well. However, the method has two major limitations. First, it allows to estimate pairwise interactions only and cannot be generalized to higher-order correlations, as it requires to define a reference and a target neuron. Higher-order correlations among $N \geq 3$ neurons may yield significant pairwise correlations among all or most neuron pairs (although this is not necessarily the case). Reversing the argument, all-to-all pairwise correlations may be an indication of higher-order interactions. Grouping pairs of all-to-all significantly cross-correlated neurons into *cliques*, and merging together overlapping cliques, Berger et al. (2007) suggested the presence of groups of neurons clustered in space in cat visual cortex that synchronize their activity in correspondence with lightening of the visual field. However, one may not conclude of the genuine higher-order nature of these interactions on the basis of pairwise analysis only. The second critical aspect of the cross-correlation analysis is that it requires fairly long data stretches to compute a reliable estimate of the histogram. This has two drawbacks: it requires correlations to be stationary over time in

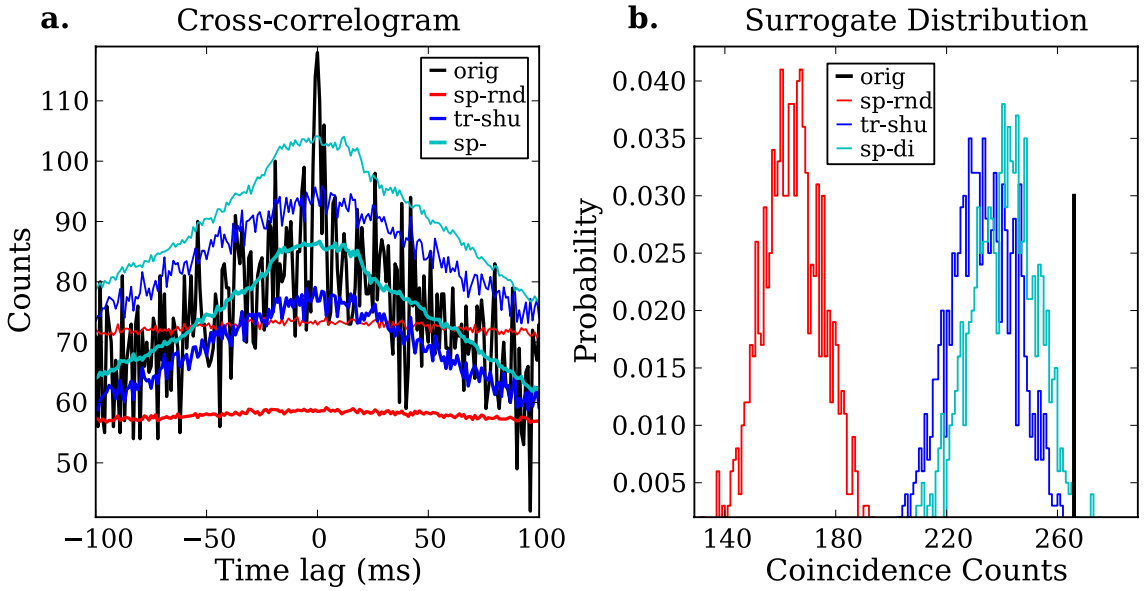

Figure 1.3: **CCH and statistical significance**. **(a)** The black line shows the row CCH resulting from the spiking activity of two fictional neurons, obtained via stochastic simulations. 40 realizations (trials) are generated for each neuron from Gamma processes ($\gamma = 3$, coefficient of variation $1/\gamma = 0.33$) having rate profiles that vary over time and across trials. In addition, jittered synchrony (± 1 ms) is injected in the data at a rate of 2 Hz. The CCH is obtained by averaging over the 40 pairs of trials. The colored lines represent the average CCH (thick lines) and the average plus twice the standard deviation (thin lines) estimated with a Monte-Carlo approach using different surrogate generation techniques: spike randomization (red), trial shuffling (blue), spike dithering (cyan). The synchrony injected in the data is responsible for the short central peak of the raw CCH, which sits on top of a broader peak caused by the ISI regularity and firing rate non-stationarities of the underlying data. Spike dithering captures these aspects and correctly predicts the broad peak, while the other methods fail in doing so. **(b)** The black vertical line shows the number of jittered synchronous events in the original data (black vertical line), computed as the integral of the central peak of the corresponding CCH in (a) from -1 ms to $+1$ ms. The colored curves represent the probability distributions of the coincidence count estimated over 1000 surrogates using different surrogate generation techniques. Figure reproduced from Grün and Rotter (2010), Figure 17.4. More details can be found in the original reference.

order to obtain reliable, unbiased estimates thereof, and prevents the analysis from being performed in a time-resolved manner, necessary to find transient, short-lasting correlations.

1.3.2 Unitary Event analysis and NeuroXidence

The Unitary Event (UE) analysis (Grün et al., 1999a, 2002a,b) was developed to overcome the limitations of the CCH analysis mentioned above, namely the limited applicability to pairwise correlations only and the impossibility to resolve the analysis in time. The analysis discretizes time into short bins (few milliseconds) and binarizes the data by representing the state of each neuron in each bin as a 1 if the neuron emits a spike in that bin, as a 0 otherwise. The network state at each bin is therefore represented by a binary vector where each element represents the state of one neuron, as illustrated in Figure 1.4A. Selected

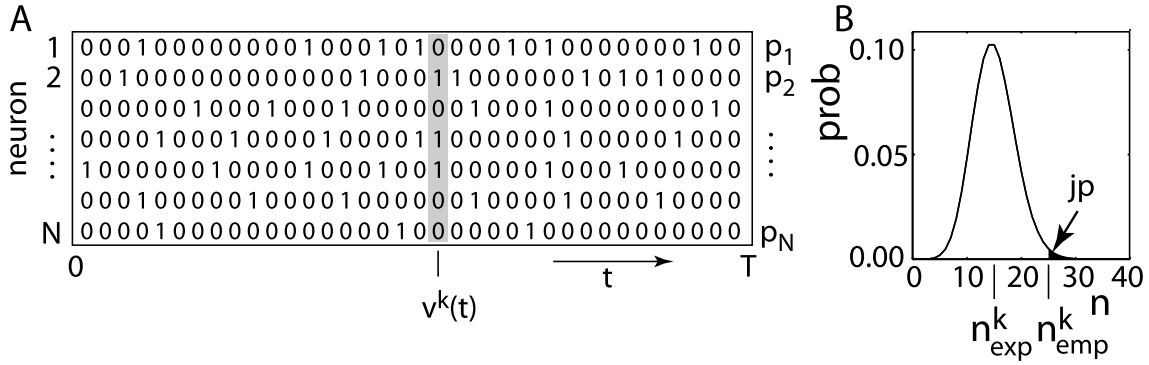

Figure 1.4: **Binary representation of spiking activity from parallel spike trains. (A)** N neurons (vertical axis) are observed simultaneously. Time (horizontal axis) is discretized into short adjacent bins, and neuronal activity is represented as a 1 (spike in the bin) or 0 (no spike). The state of the network in each bin in thus represented by a binary vector (gray-shaded area). **(B)** Computing the null distribution of the number of occurrences of each state allows mapping the empirical occurrence count into a joint probability (black area under the probability density function). Figure modified from Grün and Rotter (2010), Figure 10.1.

a time window, the analysis computes the number of occurrences of each state within that window (Figure 1.4B) and evaluates its statistical significance, either analytically or numerically by means of surrogate data. Significant states (e.g. $(1, 0, 1, 1)$) are termed Unitary Events and represent groups of neurons (e.g. neurons 1, 3 and 4) which synchronize their activity beyond the chance level.

In data consisting of a small number of neurons (and therefore taking a relatively small number of possible states), a window of about 100 ms typically provides samples of sufficient size to yield reliable estimates of the state p-values. Sliding the window over time (for instance in steps of 1 ms) allows a time-resolved analysis which captures transient UEs. Riehle et al. (1997) applied the UE analysis in a time-resoved fashion to motor-cortex data obtained from behaving monkeys during a grasping task. By doing so, they were able to find pairs and triplets of neurons (out of 6 neurons recorded in total) that synchronized their activity at times locked to behavioral events, such as visual cues that appeared or were merely expected. Kilavik et al. (2009) confirmed this finding, additionally showing that the timing of unitary events could be re-learnt after training the monkey with different waiting times.

Pipa et al. (2008) proposed a method similar to the UE analysis, called NeuroXidence, to detect excess or deficit of synchrony. The salient difference from the UE analysis is that NeuroXidence bins the data on a shorter time scale (< 1 ms) than the temporal width defining (jittered) synchronous spikes (e.g. 5 ms). This procedure mildens the performance loss due to binning in the presence of jittered synchrony. Besides, Pipa et al. (2008) proposed a new surrogate generation technique, named spike train shifting, to evaluate the null distribution. Spike train shifting generates a surrogate spike train by shifting each spike in the original train by the same amount. By doing so, it preserves not only the firing rates (which change over a slower time scale than the shift amplitude) but the autocorrelation of the spike train as well, alterations of which are a strong generator of FPs (Pipa et al., 2013b).

Despite being in principle applicable to test the significance of the occurrence count of any state of the network (i.e. any group of neurons in the sampled population), the

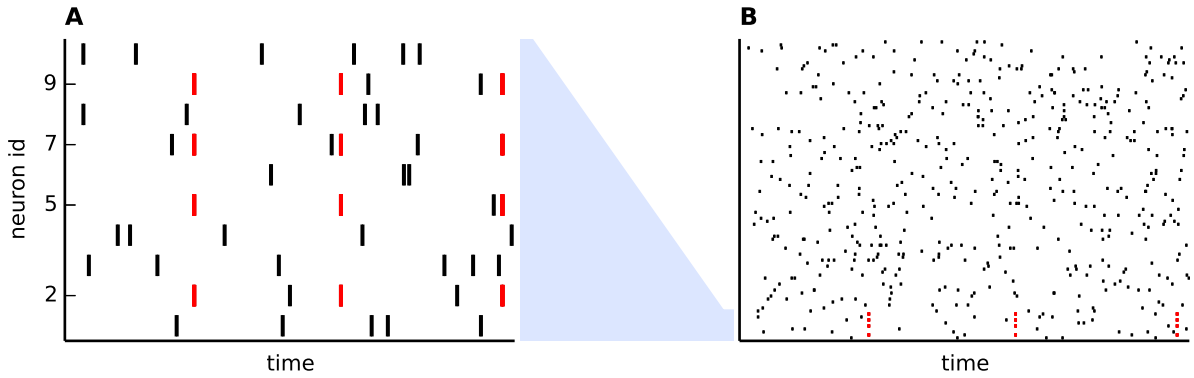

Figure 1.5: **Synchronous spike pattern embedded in low- and high-dimensional data. (A)** Cartoon raster plot of the spiking activity of 10 neurons (vertical axis) over time (horizontal axis). Each tick represents one spike. The red ticks indicate three successive occurrences of a synchrous spike event involving neurons 2, 5, 7 and 9. **(B)** The activity represented in panel (A) is embedded in a (fictional) larger network comprising 90 additional neurons. Examining a larger population has to consequences on statistical analyses of pattern significance. First, the number of possible patterns to analyze increases exponentially, raising combinatorical and multiple testing problems. Second, the probability that, when a synchronous cell assembly is active, additional neurons fire by chance increases until becoming eventually an almost sure event. In this scenario, the activity of the cell assembly is not mapped uniquely to the binary population vector representing the state (active-inactive) of all neurons.

UE analysis as well as NeuroXidence face in practice two problems when the size N of the population grows. First, the number of different states (neuron groups) that the population can form grows exponentially with the number of neurons. This number, despite being actually bounded by the number of time bins (typically some thousands for recordings lasting a few seconds), is large enough to yield severe multiple testing problems, namely high levels of false positives (FPs), or false negatives (FNs) when adopting statistical corrections. The second problem lies in the definition of a state, which is determined not only by the active neurons (1s) in the bin, but also by the silent ones ($0s$). If a group of synchronous neurons is small compared to the size of the observed population, it is likely that, when the group becomes active, a random selection of additional neurons also fires in the same time bin by chance. These neurons are likely to change at each instance of the group's activation, yielding effectively different network's states each time. When this happen, the repeated activation of the synchronous group does not boost the occurrence count of the associated state, but rather increases by a small amount the occurrence count, and therefore the p-value, of a large number of different states containing additional 1s. As a consequence, the real group of correlated neurons is very unlikely to be detected. This problem is aggravated as the relative size of synchronous cell assembly shrinks compared to that of the observed population, because more neurons external to the assembly may become active by chance at the same time of the assembly (see Figure 1.5).

1.3.3 Maximum entropy models

As seen above, the UE analysis and NeuroXidence determine the p-value of the observed states of the network by comparing their occurrence frequency in the original data with the

frequency expected under the null hypothesis of independence. One may want to relax this hypothesis, for instance by accounting for the observed pairwise correlations and testing for higher-order correlations that cannot be explained on the basis of the observed firing rates and two-ways interactions.

A multivariate distribution among N binary states (on/off neurons) has 2^N degrees of freedom, and is fully specified if all multivariate moments up to the moment of order N are fixed. When not all constraints are specified, an infinite number of different distributions exists in the general case satisfying those constraints (e.g. with given first and second moments). Each carries implicit information additional to that imposed by the contraints. When formulating a null hypothesis, one would like to minimize this additional information to make sure that no additional (and usually uncontrolled) assumption is implicitly made about the data beyond the desired constraints.

In statistical terms, minimizing the information carried by a probability distribution amounts to maximizing its Shannon entropy (Shannon, 1948). The problem of determining the least informative null distribution satisfying given constraints thus becomes a constrained maximization problem. If the constraints are averages of the associated random variables or of transformations thereof (such as the moments of the distribution), and if said constraints are consistent (i.e. allow for at least one solution), it can be proved that the solution is unique and is a Gibbs (or Boltzmann) distribution, whose expression can be derived analytically as a function of the quantities that have been constrained (Jaynes, 1957) (for details, see Appendix A.2).

Inspired by previous work (Martignon et al., 1995, 2000; Nakahara and Amari, 2002; Gütig et al., 2003), Shlens et al. (2006) and Schneidman et al. (2006) employed a maximum entropy approach to derive the null distribution of binary states (defined by means of time binning as illustrated in Figure 1.4A) in data from small numbers of spiking neurons (7 and 15 cells, respectively) in the retina. Their null hypothesis accounted for the observed average firing rates and average pairwise correlations. By doing so, both groups concluded on the absence of higher-order correlations in their data, a result which Schneidman et al. (2006) speculated could be generalized by interpolation to larger populations of up to 200 retina cells. The group proposed an index to quantify the portion of overall correlation in the data explained by a certain model, and found in their data values higher than 99% for the maximum entropy model of order 2. Qualitatively similar findings were later confirmed in further studies on the retina (Tkacik et al., 2006) and on cortical regions, such as the motor cortex (Tang et al., 2008), using recordings of similar size. However, later theoretical work by Roudi et al. (2009) demonstrated that this size is below the critical point that would allow an extrapolation of the results to populations of hundreds of neurons. Larger data have to be examined in their entirety. In Appendix A we further show on test data containing injected synchrony that values close to 1 of the index proposed by Schneidman et al. (2006) do not allow to conclude on the absence, or irrelevance, of higher-order correlations, even in data of small size.

Irrespective of this problem, the increase in data dimensionality poses two obstacles to the use of maximum entropy models. First, maximum entropy models estimate, just like the Unitary Event analysis or NeuroXidence, the distribution of specific network states as represented by binary vectors of yes/no activity. As the size of the network grows with respect to the size of an individual cell assembly therein, the activation of the latter does no longer correspond to an increased occurrence count (and therefore an enhanced statistical significance) for a single state. Second, the maximum entropy distribution of N-dimensional states constrained on its first and second moments i) has $N+\binom{N}{2} = N(N+1)/2$ parameters to be estimated, a number which grows quadratically with N, and ii) can generate 2^N different states for which to estimate the statistical significance. These aspects

raise two issues in applications to MPST data, where N is large. First, the computational effort required to estimate such a high-dimensional distribution is severe. The problem is aggravated by the fact that the expression of the Gibbs ditribution is not available in closed-form in the general case, and must be approximated numerically. This can be done by general-purpose algorithms such as iterative scaling (Darroch and Ratcliff, 1972), or by more efficient algorithms which exploit the sparsity of spikes (Broderick et al., 2007). Nevertheless, the estimation remains an open computational challenge for data cosisting of hundreds of parallel spike trains. Second, larger data also require larger samples for a correct estimation of the distribution parameters and of the frequency of observed states. This becomes a problem in applications to neuronal data, where firing rates typically vary over time and correlations may change as well. Indeed, the null hypothesis formulated by the maximum entropy approach of Schneidman et al. (2006) incorporates averages of the first and second moments of the distribution while neglecting their temporal modulation, which may bias the statistics significantly.

Kass et al. (2011), Kelly and Kass (2012) and Shimazaki et al. (2012) extended the maximum entropy framework described above to define models that capture time-varying firing rates and correlations, as well as history effects that make spike trains behave different from Poisson processes. This new class of models is a powerful tool to account for non-stationarities and history effects in data, but requires correspondingly more parameters than the simpler models above. While successful in analyzing correlations in non-stationary, non-Poisson data of small size, this approach has limited applicability to the MPST data we are interested in describing.

1.3.4 Cumulant-based inference of correlation

The spike synchrony analyses discussed above are effective in applications to low-dimensional data sets consisting of few parallel spike trains. However, they share limitations that prevent application to MPST data, most importantly the combinatorial explosion in the number of patterns, raising computational and statistical issues, and the increase of background activity that overlaps with existing patterns changing their neuronal composition. Staude et al. (2010b) introduced a novel method called CUBIC (CUmulant-Based Inference of Correlation) to circumvent these problems by estimating a measure of population correlation rather than looking for each individual pattern.

CUBIC avoids the combinatorical explosion of the complexity of the data by considering the amount of synchronous spikes in each time bin while disregarding its specific neuronal composition. Rather than assessing the statistical significance of specific patterns of synchronous spikes, the method is designed to determine the minimum order of population synchronization (i.e. the minimum size of non-chance synchronous events) that has to be assumed in the data to explain the observed synchrony. To this end, CUBIC builds the empirical distribution of the number of synchronous spikes observed across all time bins (the *complexity distribution* of the data), and evaluates the statistical significance of its first three cumulants (directly associated to the mean, the variance, and the skewness of the distribution) under a hierarchy of null hypotheses. First, the method assumes spike train independence. If the cumulants of the complexity distribution are in ranges expected under this hypothesis, the method stops and concludes that the minimum order of correlation to be assumed to explain the observed statistics of population synchrony is $\xi = 1$ (independence). Otherwise, it assumes that the population contains at most pairwise correlations. If the cumulants of the complexity distribution are explained (not significant) under this assumption, the method returns a minimum order of correlation $\xi = 2$ for the population activity. Otherwise, it goes on including higher and higher orders of correlations

in the data, until it finds the minimum order that explains the observed cumulants.

In its original version, CuBIC derives the null distribution working under the assumption that the spike trains are Poisson and time-stationary processes. Successive extensions account for deviations from these assumptions, such as Gamma-distributed inter-spike intervals (ISIs) or time-varying firing rates (Staude et al., 2010a; Reimer et al., 2012). Reimer et al. (2013) also proposed a modification of the method able to infer population correlations from membrane potentials. The method as such, however, infers the minimum order of correlation to be assumed in the data without searching for specific groups of neurons that synchronize their spiking activity. The underlying statistical model is a compound Poisson process (CPP), where synchronous events are present but involve random groups of neurons each time rather than specific cell assemblies. The minimum order ξ of correlation inferred by CUBIC does not allow to conclude on the size of cell assemblies in the data, or even on their very presence. Indeed, synchrony may be entirely caused by distributed network correlations that do not engage specific groups of neurons. Conversely, the presence of synchronous cell assemblies - which methods targeting individual spike patterns such as the UE analysis may reveal - may remain invisible to CUBIC, which operates on population statistics where the statistical significance of individual patterns is washed out. In conclusion, CUBIC offers the remarkable advantage to be applicable to MPST data, but has the drawback of doing so at the expense of the resolution of the found correlation structure, which is collapsed to a single number of difficult, often not unique interpretation.

1.4 New tools for the analysis of spike synchronization in large data

As discussed in the previous section, a variety of methods have been developed to analyze correlations in parallel spike trains. These methods have been successfully applied to simultaneous recordings of typically small numbers of neurons to demonstrate the emergence of synchronous spikes in relation to behavior in a variety of animal species, brain areas and behavioral conditions. Nevertheless, the small size of these data does not allow to conclude on the structure of synchronous spiking in larger populations. Indeed, the subsampling operated when recording a handful of neurons randomly picked within a cortical region is severe, and makes it extremely unlikely to observe functional assemblies as those hypothesized by Hebb (1949). Nowadays, rapid advances in electrophysiological technology to record from hundreds or thousands of neurons simultaneously (Buzsaki, 2004; Schwarz et al., 2014) are overcoming this issue. Consequently, the analysis of pairwise correlations in large neural networks, as well as of higher-order correlations, is becoming a possibility. Unfortunately, classical methods such as the CCH analysis, the UE analysis and maximum entropy models cannot be applied in a straightforward way to these high-dimensional data, because the combinatorical growth in the number of patterns to analyze raises severe computational and statistical problems. Alternative methods, such as CUBIC, analyze correlations in large data at the population level, giving up on the task of assessing whether the observed synchrony is generated by specific groups of correlated neurons, and even worse of determining such groups. Extrapolating results obtained from sub-populations to populations which are orders of magnitude larger, as proposed by Schneidman et al. (2006), is also not an option (see Roudi et al., 2009, and Appendix A).

These considerations call for new techniques enabling the study of synchronized spiking in high-dimensional data sets. These techniques have to cope with the variability of neuronal responses over time and across trials and units (Grün, 2009), and with the consequent paucity of samples at hand. Methods that can work on short stretches of data

would also allow for a time-resolved analysis of synchrony, a pre-requisite for the analysis of transient, behaviorally-locked correlations. In this thesis we introduce two novel statistical methods for the analysis of synchronous spike events in MPST data. The first method, named SPADE (Synchronous PAttern Detection and Evaluation), is designed to search for individual groups of neurons that synchronize their spiking acivity beyond the chance level expected on the basis of their firing rate profiles. In Chapter 2 we introduce SPADE and calibrate it extensively on test data obtained by stochastic simulations. The test data replicate typical properties of electrophysiological data, such as non-stationarity of the rates over time and across neurons.

In Chapter 3 we apply SPADE on simultaneous recordings obtained from two monkeys engaged in an instructed-delay reach-to-grasp task, where the animals grasp and pull an object for reward. The data are collected through a 100-electrodes Utah array implanted in motor cortex and covering parts of pre-motor and primary motor cortex. We perform the analysis separately for each of four trial types, differing by the grasping modality to use and the force level required to pull the object, and for each of six successive trial epochs from trial start to reward. This framework allows us to associate patterns to different behavioral contexts, which clearly determine some features such as the amount of significant patterns, their neuronal composition, their spatial orientation.

In Chapter 4 we introduce a second method, named ASSET (Analysis of Sequences of Synchronous EvenTs), to detect temporal sequences of synchronous spikes in MPST data. The interest in such temporal sequences is motivated by the assumption that the brain, if using synchrony as a meaningful tool of information coding, may support mechanisms of synchrony propagation to transmit information between distant areas. These considerations are supported by a rich literature demonstrating the possibility to define neural network models with biologically realistic constraints and enabling this dynamics, such as the synfire chain model (Abeles, 1991; Diesmann et al., 1999; Ikegaya et al., 2004; Hosaka et al., 2008). In this context, we show that ASSET is an efficient and reliable tool to reveal sequences of synchronous events in large data, where this concerted activity could first become apparent.

We conclude in Chapter 5 by reviewing the merits and limits of the two novel methods we propose for the analysis of spike synchrony and synchrony propagation, and by discussing novel insights delivered by the analysis of the experimental data about the spatial and temporal organization of synchronous spiking in motor cortex in relation to behavior. Future research directions are also proposed.

Chapter 2

SPADE: synchronous pattern detection and evaluation in MPST data

This chapter presents a novel technique for the analysis of pairwise and higher-order synchronous spike events in MPST. Our work builds on and refines a previous method proposed by Picado-Muiño et al. (2013), that i) employs frequent item set mining (FIM) to count pattern occurrences in an efficient way, and ii) estimates pattern significance on the basis of the pattern size (that is, the number of participating neurons) and pattern *support* (i.e number of pattern occurrences). The second step, called *pattern spectrum filtering* (PSF), is achieved by means of a Monte Carlo approach which builds the null distribution on correlation-free surrogates of the original data.

Here, we improve the method by Picado-Muiño et al. (2013) in three respects. First, we refine the statistical test employed in PSF to assess pattern significance (Section 2.1). Then, we introduce a further analysis step, termed *pattern set reduction* (PSR), to additionally filter out those patterns that are detected as significant on the basis of their size and support, but actually result from a combination between the occurrences of a "true" pattern and additional chance spikes (Section 2.2). We call the method resulting from the combination of FIM, PSF and PSR the Spike PAttern Detection and Evaluation (SPADE) analysis.

Here we investigate the performance of SPADE for various features of the data (coincidence rate, pattern size, firing rate variability over time or heterogeneity across neurons) and for various analysis parameters (Section 2.3). The discussion includes a step-by-step instruction on how to utilize the proposed method in the context of MPST data obtained from electrophysiological recordings. In Chapter 3 we apply SPADE to data from monkeys engaged in an instructed-delay reach-to-grasp task and investigate the emergence and features of synchronous spike patterns in relation to behavior.

Large part of this work has been published in Torre et al. (2013). Emiliano Torre designed the improved version of PSF, derived together with Christian Borgelt and David Picado-Muinõ the novel PSR approach, designed together with Michael Denker and Sonja Grün the statistical models used for the calibration, and performed the analyses presented in the paper.

2.1 Detection and statistical assessment of synchronous spike patterns

In this section we introduce our approach to detect frequent synchronous spike patterns in MPST. We first briefly review frequent item set mining (FIM) and related terminology and definitions as proposed in Picado-Muiño et al. (2013) as a tool to efficiently detect and count synchronous spike patterns in MPST. We then derive a modified version of the FIM-based statistics proposed in Picado-Muiño et al. (2013) for assessing pattern significance.

2.1.1 Frequent Itemset Mining

Given N parallel spike trains with neuron ids $1, 2, \ldots, N$, observed in the time window $[0, T)$, we partition $[0, T)$ into b exclusive bins $\{b_i\}_{i=1}^{b}$ of identical width $w = T/b$ (typically chosen as a few ms): $b_i = [(i-1) \cdot w,\ i \cdot w)$. If one or more spikes of one neuron fall into a bin, we consider the bin occupied and reduce the entry to 1 (*clipping*), so that the value taken by each time bin reflects the number of active neurons in there, rather than the number of spikes. Spikes from different neurons falling into the same time bin are defined as *synchronous* (see Figure 2.1A). Borrowing terminology from FIM, we define each neuron id as an *item*, the set T_i of all items spiking in b_i as the i-th *transaction* in the binned data, and $\{T_i\}_{i=1}^{n}$ as the *transaction list*. Given a *minimum pattern size* z_0, each set of $z \geq z_0$ items in T_i constitutes a *pattern of synchronous spikes*, or *item set* (see Figure 2.1B). Here, we set z_0 to 2. Due to clipping, each item set occurs at most once per transaction. The number of occurrences of an item set in the transaction list is the *support* of that item set.

A transaction that contains K items yields $2^K - K - 1$ different (but possibly overlapping) item sets of size $z \geq 2$, that is, all 2^K possible subsets without the empty set and the K singletons. The total number of different item sets in a transaction list can thus largely exceed the number of transactions (i.e. time bins). This number grows with the duration of the data set (see Figure 2.1D) and with the number of parallel spike trains (not shown).

In order to limit the data to potentially interesting and non-trivial item sets, we select only item sets whose support c is larger than or equal to a *minimum support* c_0 ($c_0 \geq 1$) as introduced by Picado-Muiño et al. (2013). Here we set c_0 to 2. An item set whose support equals or exceeds the minimum support is called *frequent item set*. For $c_0 > 1$, frequent item sets are usually a small fraction of all item sets (Figure 2.1D, compare black dashed line to blue dashed line). Furthermore, we discard any frequent item set occurring as many times as any of its supersets. These patterns are trivially explained by the occurrences of their supersets, which are more significant due to the larger number of neurons involved. Non-trivial frequent item sets are called *closed frequent item sets* (CFISs; see Figure 2.1C). Discarding non-closed frequent item sets does not yield any loss of information. Indeed, the set $\mathcal{F}$ of all frequent item sets can be reconstructed from the set $\mathcal{C}$ of CFISs by

$$\mathcal{F} = \bigcup_{I \in \mathcal{C}} \bigcup_{J \subset I,\, |J| \geq z_0} J.$$

The support $s(I)$ of a non-closed frequent item set $I \in \mathcal{F}$ can be computed as $s(I) = \max_{J \in C,\, J \supset I} s(J)$.

If A and B are two CFISs such that $B \subsetneq A$, and c_A, c_B their respective supports, it follows from the definition of CFISs that $c_B > c_A$ (*a priori* property). We refer to the (non-empty) set $A \backslash B$ as the *excess items* of A with respect to B, and to the difference $c_B - c_A$ as the *excess occurrences* of B with respect to A.

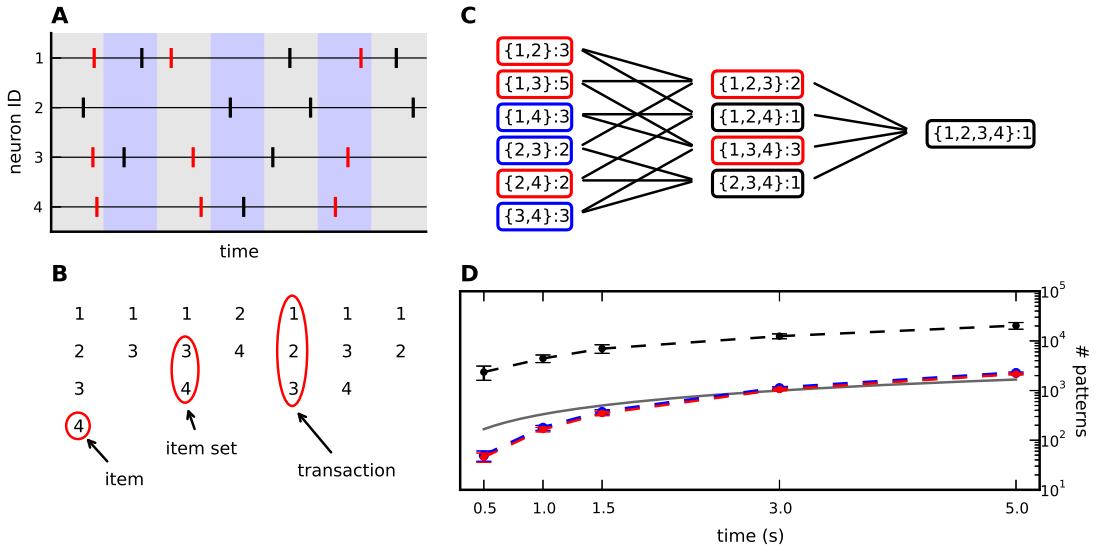

Figure 2.1: **From spike data to closed frequent itemsets**. **(A)** Sketch of a raster plot of 4 neurons firing in parallel. Shaded colors separate adjacent bins. Red spikes mark the occurrences of the synchronous pattern composed of neurons $1, 3, 4$. **(B)** Transaction list derived from the spike data in (A) after binning. **(C)** List of item sets obtained from (B), together with their occurrence counts. Black boxes mark non-frequent item sets (support set to 2), blue boxes mark non-closed frequent item sets, red boxes mark CFISs. **(D)** Average number of item sets (dashed black line), frequent item sets (dashed blue line) and CFISs (dashed red line) obtained from 100 simulations of 100 parallel independent spike trains with a firing rate of 20 Hz, as a function of the simulation time. Other parameters are bin width $w = 3$ ms and minimum pattern size $z_0 = 2$. Bars mark ± 1 std. dev. The solid line indicates the number of time bins (and thus transactions) as a function of the simulation time.

Following Picado-Muiño et al. (2013), we make use of frequent itemset mining (FIM; for a review, see Goethals, 2010; Borgelt, 2012) to extract CFISs and their support from an MPST transaction list. FIM performs a non-redundant search for spike patterns, starting from those of size z_0 and then moving on to supersets of increasing size. Starting at lowest-size patterns, the search is organized in a search tree in layers of increasing pattern size. A branch connects two patterns if one is a subset of the other. Each pattern is visited at most once. FIM exploits the *apriori* property to stop the search at infrequent patterns, as no supersets of an infrequent item set can be frequent. The output of FIM is a list of all CFISs with their support (Figure 2.1C).

2.1.2 Pattern spectrum filtering

Direct statistical tests of all individual patterns occurring in MPST data are not suitable, as they raise a severe multiple testing problem yielding large occurrences of false positives (FPs), or enhanced levels of false negatives (FNs) after statistical corrections. Therefore Picado-Muiño et al. (2013) proposed to pool CFISs according to their size z and their support c in a two-dimensional histogram (*pattern spectrum*) and to assess pattern's significance on the basis of the p-value of the pattern's signature $\langle z, c \rangle$. The signature's p-value was estimated with a Monte-Carlo approach by repeatedly generating uncorrelated surrogate data, and by evaluating the fraction of surrogates containing that signature.

Here we present a refinement of this original approach, named *pattern spectrum filtering* (PSF), that evaluates the p-value of each signature $\langle z, c\rangle$ by an upper-tail test, that is, as the fraction of surrogates containing patterns of size z *or higher*, and support c *or higher*.

Specifically, we repeatedly generate surrogate data, collect from each surrogate its CFISs by means of FIM, as done for the original data, and compute the corresponding surrogate pattern spectrum. The surrogates are generated from the original data by intentionally destroying correlations while keeping other features, such as firing rates, intact (e.g. by spike time randomization or spike dithering, see Louis et al., 2010).

Let $\succeq$ be the partial ordering on the real plane, that is, $(x_*, y_*) \succeq (x, y)$ if $x_* \geq x$ and $y_* \geq y$, where $\succ$ holds if at least one inequality is strict. From each surrogate pattern spectrum we compute a binary spectrum which takes value 1 at each signature $\langle z, c\rangle$ such that at least one signature $\langle z_*, c_*\rangle \succeq \langle z, c\rangle$ is occupied, and value 0 otherwise (in contrast to Picado-Muiño et al. (2013) where only the occupation of signature $\langle z, c\rangle$ is checked). Formally, we define the *signature operator* $\mathrm{sgt}(\cdot)$ such that, given a CFIS A with size $z_A = |A|$ and occurrence count c_A, $\mathrm{sgt}(A) := (z_A, c_A)$. For each list $\mathcal{S}_i$ of CFISs from one surrogate data set, let $\hat{P}_i$ be the *binary pattern spectrum*, defined for each $z, c \geq 2$ by:

$$\hat{P}_i(z,c) := \begin{cases} 1 \text{ if } \exists A \in \mathcal{S}_i : \mathrm{sgt}(A) \succeq (z,c) \\ 0 \text{ otherwise} \end{cases} .$$

Averaging the binary spectra at each signature, we get the *p-value spectrum* $\hat{P}$:

$$\hat{P}(z,c) := \frac{1}{K} \#\left(\mathcal{S}_i : \exists A \in \mathcal{S}_i : \mathrm{sgt}(A) \succeq (z,c)\right).$$

$\hat{P}(z, c)$ yields an estimate of the probability to observe (one or more) patterns with signature $\langle z_*, c_*\rangle \succeq \langle z, c\rangle$ under $\mathcal{H}_0$ (see Figure 2.2D).

We then classify any signature $\langle z, c\rangle$ whose p-value is lower than significance level α_* as significant. Given the desired overall significance level α for PSF, we derive α_* from α by Bonferroni correction for the number m of tests, i.e. the number of signatures in the data to test for: $\alpha_* = \alpha/m$. Any signature $\langle z, c\rangle$ for which $\hat{P}\langle z, c\rangle < \alpha_*$ is classified as significant. Formally, we introduce the *significance spectrum* $\hat{S}$ defined at each $\langle z, c\rangle$ by

$$\hat{S}\langle z,c\rangle := \begin{cases} 1 & \text{if } \langle z,c\rangle \text{ is significant} \\ 0 & \text{otherwise} \end{cases} .$$

In Figure 2.2E $\hat{S}\langle z, c\rangle = 1$ is marked in white, $\hat{S}\langle z, c\rangle = 0$ in gray. The border between the two is the *detection border*, on the left of which signatures in the original data are classified as not significant and rejected. Signatures to its right ($\hat{S}\langle z, c\rangle = 1$) are considered as significant (marked in red in Figure 2.2E). The corresponding patterns and their supports are listed in Figure 2.2F.

2.2 Pattern set reduction

PSF tests the significance of patterns under the null hypothesis H_0 of fully uncorrelated spike trains. However, PSF might fail in rejecting patterns that result from combinations of chance spikes or chance patterns with the assembly pattern (see list of detected patterns in Figure 2.2F besides the injected one). These patterns are a specific kind of false positives, not resulting from merely independent data. They may be subsets or supersets of the assembly pattern, or patterns that partially overlap with it (Figure 2.3A-C). In this section we define the type of FPs that may occur, investigate why PSF is prone to return such

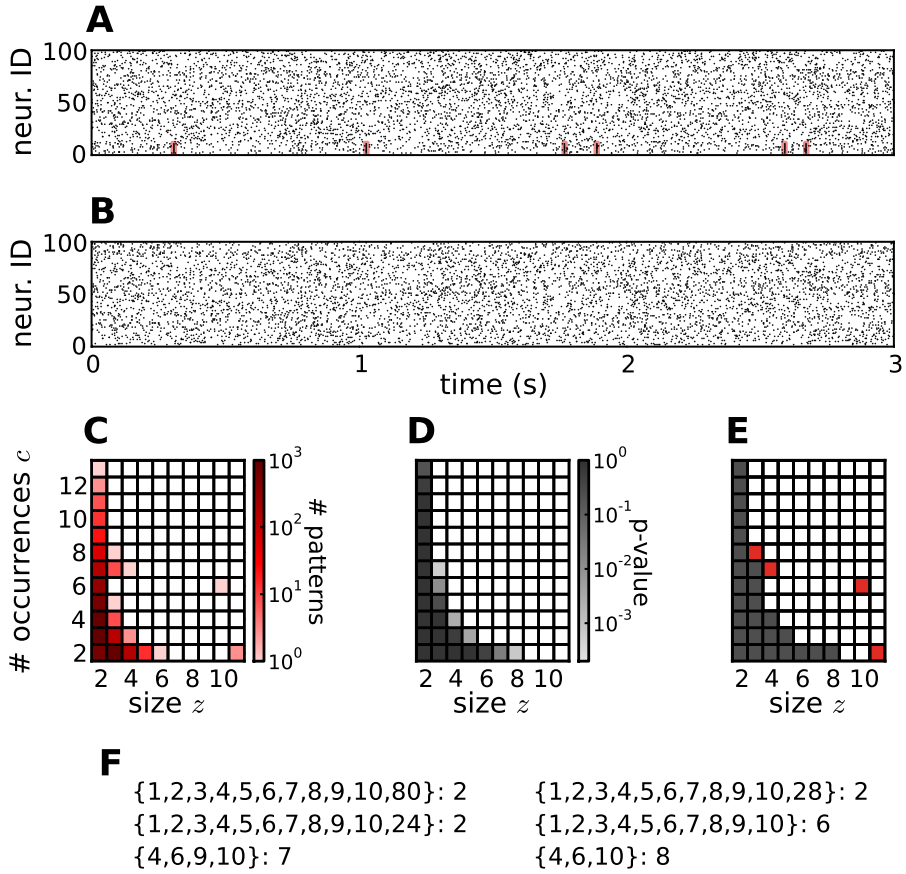

Figure 2.2: **PSF on artificial data.** **(A)** Raster plot of 100 parallel simulated spike trains consisting of independent Poisson activity plus 6 injections of one pattern of synchronous spikes (highlighted in red) from neurons 1 to 10, occurring at random times (see Section 2.3 for details). The total firing rate of each neuron is 20 Hz, the simulation time is 3 seconds. **(B)** Same as in (A), but without injection of synchronous patterns. The spike trains are therefore completely independent. **(C)** Pattern spectrum of CFISs extracted from the data in (A) by FIM ($z_0 = 2$, $c_0 = 2$, $w = 5\,\text{ms}$). Counts are color-coded (logarithmic scale). **(D)** P-value spectrum drawn from 5000 surrogate, independent data sets of the type shown in (B). P-values are color-coded (logarithmic scale). **(E)** Significance spectrum (overall significance $\alpha = 0.01$, Bonferroni-corrected for $m = 50$ tests yielding $\alpha_* = 2 \cdot 10^{-4}$). Gray squares indicate signatures that are not significant, white squares mark potentially significant signatures. Red squares mark significant signatures of the pattern spectrum shown in (C), i.e. which fall into white squares of the significance spectrum. **(F)** List of patterns detected by PSF. Besides the injected pattern $A = \{1, 2, 3, 4, 5, 6, 7, 8, 9, 10\}$, PSF also classifies additional patterns as significant, all being subsets or supersets of A.

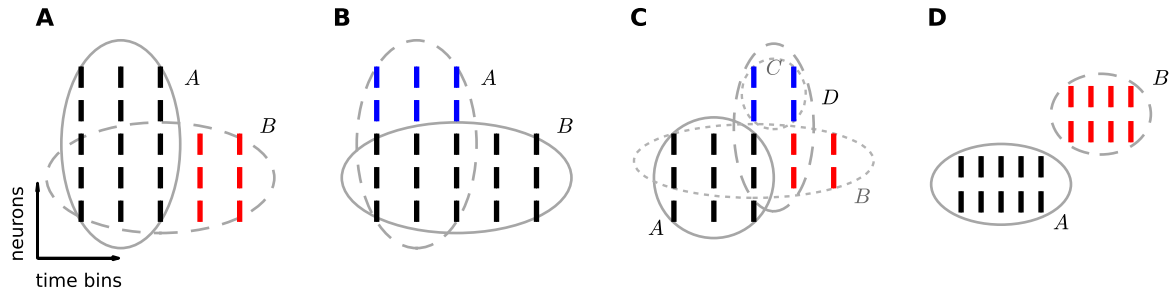

Figure 2.3: **Excess occurrences and excess items.** Sketch of the possible relationship between a reference pattern and patterns sharing neuron identities and/or time occurrences with it. In each panel, ticks represent individual spikes. Rows correspond to neurons and columns to transactions, i.e. time bins. Spikes forming a pattern are grouped by an ellipse. The reference pattern of each panel is shown by black ticks and is indicated by a solid ellipse. **(A)** B is a subset of A with excess occurrences (red). **(B)** A is a superset of B with excess items (blue). **(C)** B is a subset of A with excess occurrences (red). Neurons in C (blue) additionally fire synchronously to A and to excess occurrences of B. Thus pattern $D = B \cup C$ forms a CFIS, which partially overlaps with A. **(D)** Patterns A and B are disjoint: they are composed of different neuron identities and occur at different time bins.

FPs, and propose an additional statistical analysis, termed *pattern set reduction* (PSR), to remove them.

2.2.1 Types of FPs

Chance subsets. If a CFIS A repeats c_A times and a subset B of A (with $|B| \geq z_0$) has c additional chance occurrences, B represents a CFIS repeating $c_B = c_A + c$ total times. We call B a *chance subset* of A, having c excess occurrences (Figure 2.3A). PSF is designed to test the significance of signature ($|B|$, c_B) under H_0 (complete independence), thus disregarding the fact that c_A occurrences are due to pattern A. As a result it classifies B as a significant pattern, thus yielding an FP outcome. This is illustrated in Figure 2.2F, where e.g. pattern $\{4, 6, 10\}$ occurs twice by chance plus 6 times as a subset of pattern $\{1, 2, 3, 4, 5, 6, 7, 8, 9, 10\}$. The corresponding signature $(3, 8)$ is significant compared to the surrogates (Figure 2.2E), so that PSF does not reject it.

Chance supersets. If a CFIS B occurs c_B times and another set C of neurons fires by chance synchronously with B in c of those c_B transactions (with $c \geq c_0$), then the pattern $A = B \cup C$ represents a CFIS repeating $c_A = c$ times. We call A a *chance superset* of B, with $|C|$ excess neurons (Figure 2.3B). PSF tests the significance of signature ($|A|$, c_A) under H_0, disregarding the fact that $|B|$ of the $|A|$ neurons of A are due to the presence of pattern B. The test is therefore prone to classify A as significant. This is the case for patterns $\{1, 2, \ldots, 10, 80\}$, $\{1, 2, \ldots, 10, 28\}$ and $\{1, 2, \ldots, 10, 24\}$ in Figure 2.2F, each of which occurs twice as a superset of $\{1, 2, \ldots, 10\}$. The corresponding signature $(11, 2)$ is significant compared to the surrogates (Figure 2.2E), so that PSF classifies these patterns as significant.

Chance overlapping sets. The simultaneous presence of excess items and excess occurrences can yield yet another type of FP outcome, namely patterns that overlap with

the actual assembly. Given an assembly A, assume that a subset B of A has additional chance occurrences. If an additional set C of neurons disjoint from A fires synchronously to A *and* to an excess occurrence of B for a total of $c \geq c_0$ chance times, then the set $D = B \cup C$ represents a CFIS which partially overlaps with A (Figure 2.3C). PSF is prone to classify D as significant.

Disjoint patterns Two patterns which do not have items in common are *disjoint* (Figure 2.3D). In contrast to the previous classes of chance patterns, the presence of an active assembly does not enhance chance patterns disjoint from it. PSF therefore correctly estimates their significance and manages to filter out almost all of them, as shown in Section 2.3.

2.2.2 PSR statistics

Let $\mathcal{P}$ be the class of CFISs reported as significant by PSF. Given a pair $(A, B) \in \mathcal{P} \times \mathcal{P}$ such that $B \subset A$ (therefore $c_B > c_A$ by definition of CFIS, and $|B| < |A|$), we propose statistical tests to assess the conditional significance of either A *given* B ($A|B$) or B *given* A ($B|A$)**,** i.e. of one pattern given that the other represents an assembly pattern. These tests can be applied, using different strategies, to the class of all such (A, B) pairs, reducing $\mathcal{P}$ to a subclass $\mathcal{Q}$ of patterns which are mutually significant given each other.

Subset filtering.

This procedure aims at rejecting FPs that are chance subsets of other CFISs. For each pair $(A, B) \in \mathcal{P} \times \mathcal{P}$ such that $B \subset A$ (so that $c_B > c_A$), B has $c_B - c_A$ excess occurrences with respect to A. Subset filtering tests $B|A$, i.e. the null hypothesis $\mathcal{H}_0^{B|A}$ that B is a chance subset of the actual assembly A, by assessing the significance of the excess occurrences of B. Equivalently, $\mathcal{H}_0^{B|A}$ states that the pattern B' defined by the same items as B but its excess occurrences only (red spikes in Figure 2.3A) is a chance pattern. If $\mathcal{H}_0^{B|A}$ is rejected, B is kept and A discarded, otherwise A is kept and B discarded. Thus, the procedure keeps either A or B and discards the other (*exclusive*).

We present two alternatives to test $\mathcal{H}_0^{B|A}$.

Exact test. This test computes the p-value of the signature $(|B|, c_B - c_A)$ of B'. If $c_B - c_A < c_0$, B is classified as a chance subset of A. Otherwise, let T'_A be the transaction list obtained from T by discarding the transactions where A occurred, and keeping in the remaining transactions only the items composing A. All the excess occurrences of subsets of A must be contained in T'_A. B' itself is a CFIS in this transaction list: it is an item set because $|B'| = |B| \geq z_0$, it is frequent because $c_B - c_A \geq c_0$, it is closed because otherwise B itself would be non-closed. To test the significance of B', one can therefore run FIM and PSF on surrogates of T'_A to estimate the significance of its signature $(|B|, c_B - c_A)$**.** If $(|B|, c_B - c_A)$ is significant, B' is significant in T'_A and B is classified as significant in T (given A). Otherwise, B is classified as non-significant.

Approximate test. This test approximates the p-value of the signature $(|B|, c_B - c_A)$ in T'_A by the p-value of the signature $(|B|, c_B - c_A + h)$**,** $h \geq 1$, in T, already obtained when performing PSF. In contrast to T'_A, T is composed of more neurons than those which can actually form chance subsets of A (because it does not contain the items of A only), and more transactions than those where such subsets could actually display excess occurrences (because it also contains the transaction where A is already present). Therefore, the p-value

of $(|B|, c_B - c_A)$ would be underestimated if computed over T instead of T'_A. Parameter h heuristically corrects for this by substituting it with the p-value of a signature with the same size but higher support. The lower h, the higher the probability to reject B. If $h \geq c_A$, then $(|B|, c_B - c_A + h) \succeq (|B|, c_B)$ and B is necessarily reported as significant. This test avoids to run FIM and PSF on T'_A and is therefore computationally more efficient.

Superset filtering.

This procedure aims at rejecting FPs that are chance supersets of other CFISs. For each pair $(A, B) \in \mathcal{P} \times \mathcal{P}$ such that $B \subset A$ (so that $|B| < |A|$), A has $|A| - |B|$ excess items with respect to B. Subset filtering tests $A|B$, i.e. the null hypothesis $\mathcal{H}_0^{A|B}$ that A is a chance superset of the actual assembly B, by assessing the significance of the excess items of A. Equivalently, $\mathcal{H}_0^{A|B}$ states that the pattern A' defined by the same transactions as A but containing its excess items only (blue spikes in Figure 2.3B), is a chance pattern. If $\mathcal{H}_0^{A|B}$ is rejected, A is kept and B discarded from $\mathcal{P}$, otherwise B is kept and A discarded from $\mathcal{P}$. Thus, the procedure keeps either A or B and discards the other (*exclusive*).

We present two alternatives to test $\mathcal{H}_0^{A|B}$.

Exact test. This test computes the significance of the signature $(|A| - |B|, c_A)$ of A'. If $|A| - |B| < z_0$, A is classified as a chance superset of B. Otherwise, let $T'_{\bar{B}}$ be the transaction list obtained from T by keeping only the transaction where B occurred, and discarding from them the items constituting B. All groups of excess items of B (i.e. neurons that fire synchronously to B) must be contained in $T'_{\bar{B}}$. A' itself is a CFIS of this transaction list: it is an item set because $|A'| = |A| - |B| \geq z_0$, it is frequent because $c_A \geq c_0$, it is closed because otherwise A itself would be non-closed. To test the significance of A', one can therefore run FIM and PSF on surrogates of $T'_{\bar{B}}$ to estimate the p-value of its signature $(|A| - |B|, c_A)$**.** If $(|A| - |B|, c_A)$ is significant, A' is significant in $T'_{\bar{B}}$ and A is classified as significant in T (given B). Otherwise, A is classified as non-significant.

Approximate test. This test approximates the p-value of the signature of A' in $T'_{\bar{B}}$ by the p-value of signature $(|A| - |B| + k, c_A)$**,** $k \geq 1$, in T, already obtained when performing PSF. In contrast to $T'_{\bar{B}}$, T is composed of more neurons than those that can actually form excess items of B (because it contains the items of B too), and more transactions than those where supersets of B could actually occur (because it contains also transactions where B does not occur). Therefore, the p-value of $(|A| - |B|, c_A)$ would be underestimated if computed over T instead of $T'_{\bar{B}}$. Parameter k heuristically corrects for this by substituting it with the p-value of a signature with the same support but higher size. The lower k, the higher the probability to reject A. Note that if $k \geq |B|$ then $(|A| - |B| + k, c_A) \succeq (|A|, c_A)$ and A is necessarily reported as significant.

This test allows to avoid running FIM and PSF on $T'_{\bar{B}}$ for each B.

Covered-spikes criterion.

This simple selection strategy consists of taking all pairs $(A, B) \in \mathcal{P} \times \mathcal{P}$ for which $B \subset A$, and keeping for each pair the pattern covering the largest number of spikes, while rejecting the other. Specifically, the criterion prefers A to B if $z_A \cdot c_A \geq z_B \cdot c_B$, B to A otherwise. It does not involve significance tests, but is based on the observation that, given the probability p for a neuron to spike in a time bin, the probability for z neurons to fire synchronously in a bin is approximately p^z, so that the probability that this pattern occurrs c times is binomially distributed and approximately proportional to $p^{z \cdot c}$. The larger the

$z \cdot c$ score, the less likely a pattern of such size and support. This matches the finding that the detection border separating non-significant signatures (marked gray in Figure 2.2E) from significant ones (marked white in Figure 2.2E) in the significance spectrum exhibits a hyperbolic shape. The criterion thus keeps the less likely of the two patterns.

A variant consists in keeping the pattern with the largest $(z-1) \cdot c$ score. This choice is motivated by the observation that a pattern of size z and support c can be seen as $z-1$ spike trains which synchronize their spikes to another train c times. Thus, $(z-1) \cdot c$ spikes are coincident to spikes in another spike train. Keeping the pattern with the largest $(z-1) \cdot c$ score amounts to keeping the pattern which covers more coincident spikes. Geometrically, penalizing the pattern size corrects for the fact that the hyperbolic shape of the detection border in Figure 2.2E is elongated towards the pattern support (y-axis) rather than being equilateral.

Combined filtering.

Subset filtering, superset filtering and covered-spikes criterion can be combined into a filtering procedure which tests for both excess coicidences and excess items. Combined filtering tests for each pair $(A, B) \in \mathcal{P} \times \mathcal{P}$ both the null hypotheses $\mathcal{H}_0^{B|A}$ (i.e. that B is a chance subset of A) *and* $\mathcal{H}_0^{A|B}$ (i.e. that A is a chance superset of B). If one of the null hypotheses is rejected, the corresponding pattern is retained as significant. Thus, if both hypotheses are rejected, both patterns are retained (*inclusive*). Accepting one null hypothesis does not necessarily lead to the rejection of the corresponding pattern (in contrast to subset or superset filtering): the pattern is rejected only if the other pattern is accepted, i.e. if the other null hypothesis is rejected. If both $\mathcal{H}_0^{B|A}$ and $\mathcal{H}_0^{A|B}$ are accepted, one of the two patterns is kept based on the covered-spikes criterion.

2.3 Validation on artificial data

In this section we compare the performance (in terms of FPs and FNs) of PSF to PSF followed by PSR to illustrate the advantages yielded by the latter. For the sake of computational efficiency we employ the approximate versions of the tests for the subset and superset filtering with parameters $h = 1$ and $k = 2$, respectively. We test different types of artificial data that involve typical features of experimental data. After studying the general behavior of the analysis method for stationary, homogeneous data, we study data sets with heterogeneous firing rates across neurons, and with non-stationary firing rates in time.

Correlated data. As a model for data containing assembly activity, we generate correlated spike trains by a modified version of the single-interaction-process (SIP; see Kuhn et al., 2003a; Berger et al., 2010), which we keep calling SIP for convenience. First, we simulate $N = 100$ parallel independent Poisson spike trains as background activity. Then we model assembly activity by inserting synchronous spike events in a subset of z of the N neurons (the *SIP neurons*, with ids 1 to z). This is done by generating a hidden Poisson process with the desired number c of pattern occurrences, from which spikes are copied into each of the z spike trains of the SIP neurons. Thus, as compared to the model proposed by Kuhn et al. (2003a) we insert correlated firing only in a specific subset of the parallel processes. Before insertion of the synchronous patterns, the background firing rate of the SIP neurons is reduced by the rate of the hidden process to ensure the same firing rate for all neurons. In the simplest scenario, the firing rates and the pattern occurrence rate are

Simulation parameters		Analysis parameters	
background activity	SIP	FIM	statistical tests
$N = 100$	$z = 2 - 10$	$w = 3,5$ ms	$\alpha_* = 2 \cdot 10^{-4}$
$r = 5, 10, 20$ Hz	$c = 2 - 10$	$c_0 = 2$	$K = 5000$
$T = 3$ sec		$z_0 = 2$	$R = 1000$

Table 2.1: **Parameters for calibration.** *Parameters for the background activity:* N: number of neurons, r: firing rate, T: simulation time. *Parameters for correlated data:* z: number of neurons in correlated activity (size of SIP), c: SIP occurrences. *Analysis parameters*: w: bin width, c_0: minimum item set support, z_0: minimum item set size. *Statistical parameters:* α_* : Bonferroni-corrected significance level (for $m = 50$ tests), K: number of surrogates, R: number of simulation runs per SIP model.

stationary over time and homogeneous across neurons. More complicated cases will include either non-stationarity or heterogeneity of rates. The purpose of the analysis of such data is to test under controlled conditions if the simulated assembly is indeed detected and can be distinguished from background activity.

Independent data. To implement the null hypothesis H_0 of mutual independence needed to derive the significance of signatures of the correlated data, we generate independent Poisson processes of the same rates as the data to be tested, thus keeping the same marginal statistics. This is one way of implementing the null-hypothesis. However, in the context of analyzing real experimental data, one may want to keep more statistical features of the experimental data (e.g. non-stationary and heterogenous firing rates, deviation from Poisson, and so on). This can be realized by the use of more complex surrogates derived by manipulation of the original data, e.g. spike dithering (Louis et al., 2010; Grün, 2009).

Assessing significance. We evaluate the performance of our analysis in terms of the average number of FPs and FNs obtained with PSF and PSR in $R = 1000$ iterations on the same model of correlated data (SIP of size z in $N = 100$ parallel spike trains). To study the performance of our analysis, we investigate 243 models differing in the size of the injected assembly $z = 2 - 10$, its injection count $c = 2 - 10$, and the firing rates $r = 5$, 10 or 20 Hz (here: homogeneous for all neurons). We analyse each model with a bin width $w = 3$ ms and $w = 5$ ms for the detection of synchronous spike patterns. See Table 2.1 for an overview of the parameter combinations. For the significance estimation we generate surrogate data, i.e. independent Poisson processes with the same firing rates as the correlated data, and analyse them with FIM as done for the correlated data. This procedure is repeated for $K = 5000$ times to derive the p-value and then the significance spectrum by employing an overall significance level of $\alpha = 0.01$, Bonferroni-corrected for the number of signatures tested. The latter is given by the number of signatures existent in the correlated data, which never exceeded $m = 50$. In order to have the same corrected significance level for each of the 1000 iterations of each SIP model, we always correct for $m = 50$ tests, instead of correcting for the individual number $m' < m$ of signatures found in each data set. This yields the corrected significance level $\alpha_* = 2 \cdot 10^{-4}$, which is typically more conservative than correcting individually for m' tests. This procedure allows us to use a single significance spectrum for all 81 SIP models with the same firing rates, differing by parameters z and c only, and for all 1000 realizations of each model. To obtain the p-values with precision α_* we generate $K = 1/\alpha_* = 5.000$ surrogates, compute their binary

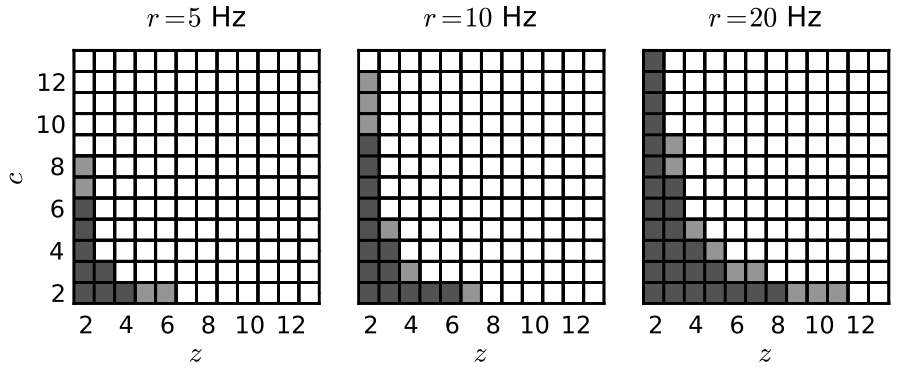

Figure 2.4: **Significance spectra for different parameter sets.** Independent Poisson spike trains ($N = 100$; $T = 3\,\mathrm{s}$) of different firing rates ($r = 5$, 10 or 20 Hz) serve as surrogates for the computation of three significance spectra (from left to right). Each square represents a $\langle z, c\rangle$ signature. Dark-shaded gray squares mark non-significant signatures obtained with $w = 3\,\mathrm{ms}$. Light-shaded squares represent further non-significant signatures for $w = 5\,\mathrm{ms}$. White squares indicate significant signatures for both choices of the bin width. Other parameters: $z_0 = 2$, $c_0 = 2$, $\alpha_* = 2 \cdot 10^{-4}$, $K = 5000$.

spectra and average them to draw the p-value spectrum (see Section 2.1.2).

Figure 2.4 shows significance spectra obtained from surrogate data for models differing by the firing rate r (5, 10 or 20 Hz) analysed with different bin widths w (dark gray for $w = 3\,\mathrm{ms}$, light gray for $w = 5\,\mathrm{ms}$; $\alpha_* = 2 \cdot 10^{-4}$). The set of non-significant signatures shows a hyperbolic shape, which grows with both r and w to higher z and higher c. Both factors, higher firing rate or larger bin width, cause more spikes per bin, and therefore larger and more frequent chance patterns.

2.3.1 Time-stationary, cross-neuron homogeneous firing rates

For each SIP parameter set we simulate the corresponding model $R = 1000$ times, and evaluate FPs and FNs of each realization. Their averages measure the performance of the analysis for each parameter constellation.

As previously discussed (Section 2.2), in the presence of correlations PSF tends to classify chance subsets, supersets or overlapping sets as significant, thus yielding FPs. Figure 2.5, top row, shows this effect on simulations of SIP models differing by SIP size (x-axis of each panel) and injection count (y-axis). For each model, the FP level is computed as an average over 1000 stochastic simulations. The total amount of FPs increases as the SIP size and/or the number of injections get larger. The contribution of FP supersets (green) and FP subsets (blue) is about the same, while in comparison FP overlapping sets (yellow) occur only at higher values for z and c, and FP disjoint patterns (purple) are almost never observed. As shown in Figure 2.5, bottom row, PSR (here, combined filtering) largely reduces the amount of FPs. Although the PSR statistical tests apply to chance subsets (blue) and supersets (green) only (Section 2.2.2), they successfully remove most of the overlapping patterns (yellow) as well. The reason is that, if there is a CFIS D overlapping with the actual assembly A by z_0 or more items, their intersection B is a CFIS as well (Figure 2.3C). In most cases PSF classifies B as significant together with A and D. If so, PSR likely rejects D when testing $H_0^{D|B}$, and rejects B when testing $H_0^{B|A}$.

A reduction of the amount of FPs typically comes at the expense of enhanced FNs. In particular, FNs may occur if the real pattern is rejected in favor of one of its subsets or supersets. Figure 2.6 shows, for a range of combinations of SIP size and injection

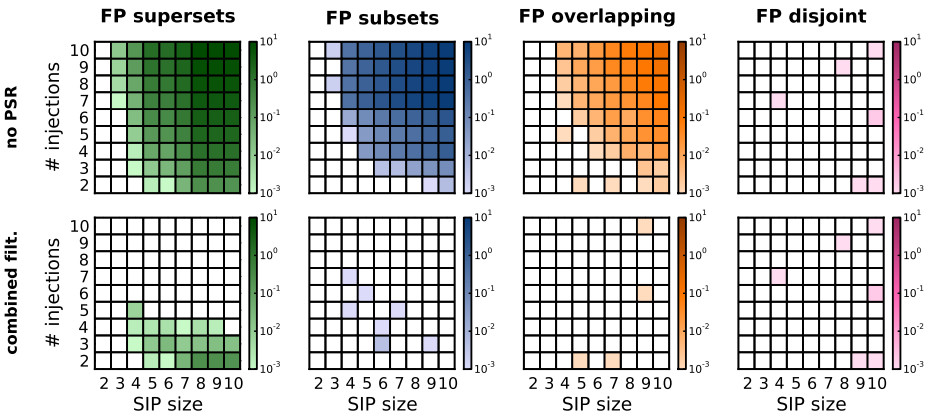

Figure 2.5: **Average number of FPs, distinguished by type, after PSF and PSR.** Average number of FPs obtained for different SIP models on $R = 1000$ model simulations. FPs are shown after performing PSF (top) and then PSR with combined filtering (bottom), and are distinguished by type (columns from left to right: FP supersets, FP subsets, FP overlapping, FP disjoint patterns). Each panel shows the average number of FPs obtained for different SIP models, each corresponding to a square in the grid: the models differ by the SIP size (from 2 to 10; x-axis) and its injection count (from 2 to 10; y-axis). Other parameters (same for all simulations): $N = 100$, $T = 3\,\mathrm{s}$, $r = 20\,\mathrm{Hz}$, $w = 3\,\mathrm{ms}$, $K = 5000$, $\alpha_* = 2 \cdot 10^{-4}$.

count, the resulting level of FPs, FNs, and the maximum of the two (as a measure of overall performance) after performing each of the proposed PSR strategies. The significance spectrum used to determine significance for all realizations of the SIP models is the one for $w = 3\,\mathrm{ms}$ shown in Figure 2.4 (top right, dark-shaded entries). For the FPs shown in Figure 2.6, top row, the color-coded level refers to the fraction of simulations (out of 1000) containing one or more FPs. This measure takes values between 0 and 1, unlike the average FP counts shown in Figure 2.5. This representation simplifies the comparison with the average FN level, which ranges between 0 to 1 since here only a single spike pattern is injected in every simulation. To aid the comparison between the performances of PSF and PSR, gray dots mark those squares that correspond to models where the error rates exceed 5%. PSF on its own never performs well in terms of FNs and FPs simultaneously, while all PSR strategies yield a range of models for which both quantities are low. In summary, the relative improvement of PSR versus PSF shows that any PSR strategy reduces the FP rate considerably, while causing only a minor increase in the FN rate.

2.3.2 Heterogeneous firing rates across neurons

If neurons have the same spiking statistics, the spike pattern statistics depends on the pattern size only. Thus, the p-value of each pattern is fully determined by the pattern signature. This does not hold when neurons have different spiking statistics, and in particular different firing rates. Here we discuss the case of heterogeneous firing rates across neurons, which are often present in electrophysiological data. Higher firing rates lead to a higher spiking probability per time bin. Patterns composed of neurons with higher firing rate are more likely to occur by chance, and are therefore less significant than patterns composed of neurons with lower rates. Thus, the p-values of patterns with the same signature $\langle z, c \rangle$ differ for different compositions of the firing rates. Pooling patterns by size and support in the pattern spectrum does not take into account the heterogeneity of firing rates across

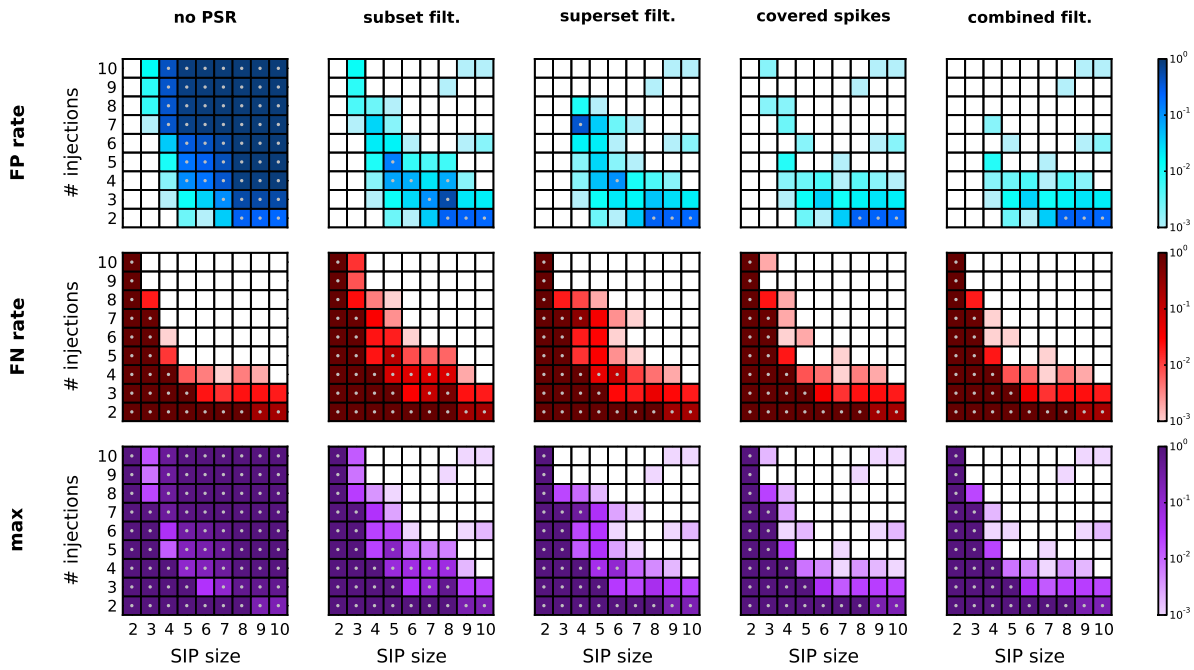

Figure 2.6: **Performance of PSR with homogeneous, stationary firing rates.** Performance of PSR with different filtering methods, measured as the fraction of $R = 1000$ simulations where FPs (top row) and FNs (second row) are detected (thus the fraction represents a rate). The maximum of the two (third row) indicates the combined error rate. Each matrix shows the performance for 81 different SIP models varying by SIP size (from 2 to 10, x-axis) and number of SIP injections (from 2 to 10, y-axis), of stationary and homogeneous neuronal firing rates ($r = 20\,\mathrm{Hz}$). The performance value is color-coded (see color bar, logarithmic scale). White squares mark SIP models where no simulations led to false outcomes. Gray dots mark entries where the error rate is above 5%. Each column corresponds to a different PSR strategy applied after PSF, from left to right: no filtering, subset filtering, superset filtering, covered-spikes criterion, combined filtering. Bar graphs (bottom) summarize the FP (red bars) and FN (blue bars) rates shown above: the upper graphs show the average FP and FN rates for each PSR method, while the lower graphs show the relative improvement over PSF. Other parameters (same for all panels): $N = 100$, $T = 3\,\mathrm{s}$, $w = 3\,\mathrm{ms}$, $K = 5.000$, $\alpha_* = 2 \cdot 10^{-4}$.

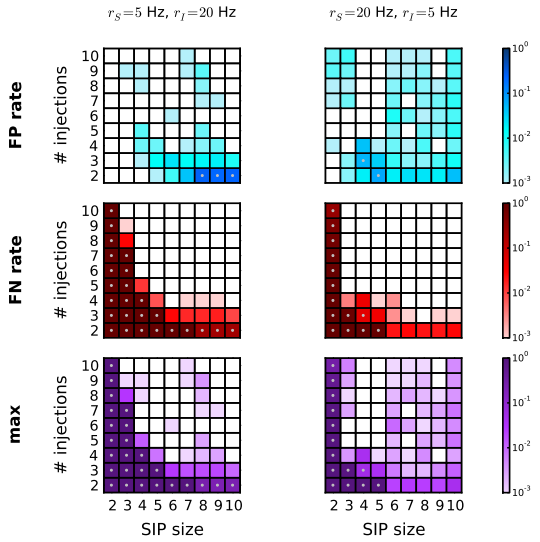

Figure 2.7: **Performance of PSR with heterogeneous firing rates.** Performance of PSR (combined filtering with parameters $h = 1$, $k = 2$) in terms of FP rates (*top* row), FN rates (middle row) and the combined error rates (maximum of FP and FN rates) (*bottom* row) of data with heterogeneous rates. Left column: SIP neurons fire at $r_S = 5\,\text{Hz}$, independent neurons fire at $r_I = 20\,\text{Hz}$. Gray dots mark entries where the error rate is above 5%. Right column: SIP neurons fire at $r_S = 20\,\text{Hz}$ and independent neurons at $r_I = 5\,\text{Hz}$. Other parameters (same for all panels): $N = 100$, $T = 3\,\text{s}$, $w = 3\,\text{ms}$, $K = 5.000$, $\alpha_* = 2{\cdot}10^{-4}$.

neurons and thus may lead to a biased statistics.

To investigate the robustness of our method against firing rate heterogeneity, we first simulate independent data consisting of 100 neurons, with a small population of neurons (2 to 10) firing at a higher rate (20 Hz) than the rest of the neurons (5 Hz). We simulate 1000 data sets of this type, and evaluate FPs in each of them by means of FIM and PSF ($K = 5000$ surrogates). In none of the simulations we detect significant signatures, i.e. FPs. The opposite scenario, where 2 to 10 neurons fire at 5 Hz and the others at 20 Hz, does not yield FPs as well. Thus, employing rate-preserving surrogates allows PSF to correctly estimate the significance of signatures under H_0, also when rates are heterogeneous across neurons.

Next we study correlated data characterized by heterogeneous background firing rates. We investigate two cases based on a SIP model. In scenario S1, a pattern is injected in a set of neurons firing with lower firing rate ($r_S = 5\,\text{Hz}$) than the independent neurons firing at rate $r_I = 20\,\text{Hz}$ (Figure 2.7, left column). In contrast, in scenario S2 the pattern is injected in neurons with higher firing rates ($r_S = 20\,\text{Hz}$, $r_I = 5\,\text{Hz}$; Figure 2.7, right column). In comparison to the homogeneous case where all neurons fire at 5 Hz (data not shown), the overall performance drops significantly, but does not so compared to the 20 Hz homogeneous case (see Figure 2.6, right column). This is consistent with the previous finding that higher rates worsen the performance by shifting the detection border in the significance spectrum to the right (Figure 2.4, left vs right). This also explains why FP and FN rates in scenario S1 are higher than in scenario S2: the average firing rate in the former ranges (depending on the SIP model) from 18.5 Hz to 19.7 Hz, in the latter from 5.3 Hz to 7 Hz. Our choice of using PSR with combined filtering leads to a better performance in this scenario than the covered spikes criterion (not shown). Taken together, these results indicate that the method can deal well with heterogeneity of firing rates without severe performance loss.

2.3.3 Time-varying firing rates

Now we want to consider the case when the firing rates of the neurons are not stationary in time. To explore the sensitivity of our method to non-stationarities we employ simulated data, again consisting of 100 parallel spike trains, which fire in two consecutive epochs of length T_1 and T_2 (the total simulation time $T = T_1 + T_2$ is 3 s, as in the data previously

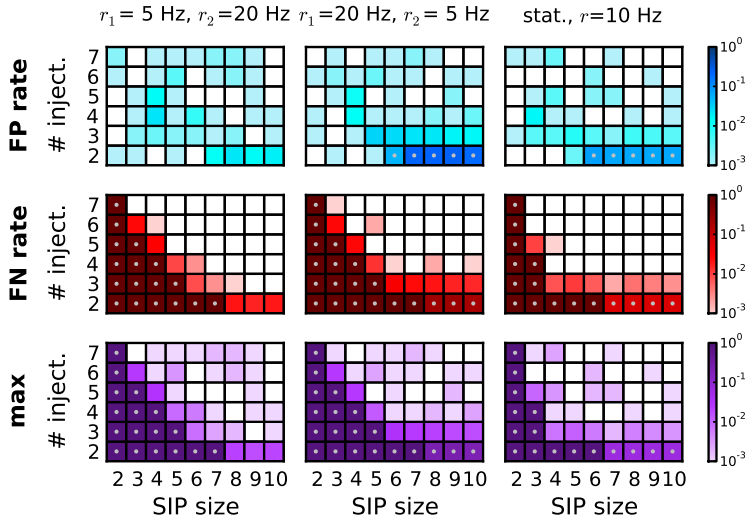

Figure 2.8: **Performance of PSR with non-stationary firing rates.** Performance of PSR (combined filtering) in terms of FP rate (top row), FN rate (middle row), and maximum of the two (bottom row), computed over $R = 1000$ simulations per SIP model. Each panel shows the performance for 54 different models varying by SIP size (from 2 to 10, x-axis) and SIP injections (from 2 to 7, y-axis). In the first two columns, the simulations consist of two epochs. The first epoch of duration $T_1 = 1.5\,\mathrm{s}$ is composed of a stationary, homogeneous SIP model (with firing rate $r = r_1$), followed by a second epoch of $T_2 = 1.5\,\mathrm{s}$ with independent spiking at a rate of $r_2 \neq r1$. In the first column, the rate compositions are $r_1 = 5\,\mathrm{Hz}$ and $r_2 = 20\,\mathrm{Hz}$, and in the second column $r_1 = 20\,\mathrm{Hz}$ and $r_2 = 5\,\mathrm{Hz}$. For comparison, the third column shows the performance for the stationary case with all neurons firing at rate $r = 10\,\mathrm{Hz}$, and a duration of $T = 3\,\mathrm{s}$. Other parameters (same for all panels): $N = 100$, $w = 3\,\mathrm{ms}$, $K = 5.000$, $\alpha_* = 2 \cdot 10^{-4}$. Gray dots mark entries where the error rate is above 5%.

analysed) at different rates ($r_1 = 5\,\mathrm{Hz}$ and $r_2 = 20\,\mathrm{Hz}$; or vice versa), homogeneously across the neurons in both epochs. In the first epoch, correlated activity in inserted by the SIP model. SIP of size 2 to 10, injected 2 to 7 times, amounting to a coincidence rate of 0.66 to 3.33 Hz in the first epoch. The background rate is reduced correspondingly. For comparison, we also study the stationary case, where all neurons fire at $r = 10\,\mathrm{Hz}$. The performance for the three scenarios is shown in Figure 2.8 (first column: $r_1 = 5\,\mathrm{Hz}$, $r_2 = 20\,\mathrm{Hz}$; second column: $r_1 = 20\,\mathrm{Hz}$, $r_2 = 5\,\mathrm{Hz}$; third column: $r_{1,2} = 10\,\mathrm{Hz}$). Although our analysis performs better (detection border more to the left) in the stationary case ($r = 10\,\mathrm{Hz}$; third column), it can still recover SIP activity with no FPs in a large portion of the parameter space, provided that rate-preserving surrogates are employed. As in the heterogeneous case, FPs increase when the SIP neurons have higher firing rates and thus more FP subsets occur. As apparent from Figure 2.8, bottom row, the method can correctly detect significant patterns in a wide range of models also in the presence of non-stationary rates. To study whether short transients in the firing rates tend to generate FPs, we repeated the analysis for $T_1 = 0.5\,\mathrm{s}$, $T_2 = 2.5\,\mathrm{s}$, setting first $r_1 = 5\,\mathrm{Hz}$, $r_2 = 20\,\mathrm{Hz}$ and then $r_1 = 20\,\mathrm{Hz}$, $r_2 = 5\,\mathrm{Hz}$. In all cases we do not find enhanced FPs (data not shown), indicating that employing rate-preserving surrogates suffices to correct for rate non-stationarity in independent data.

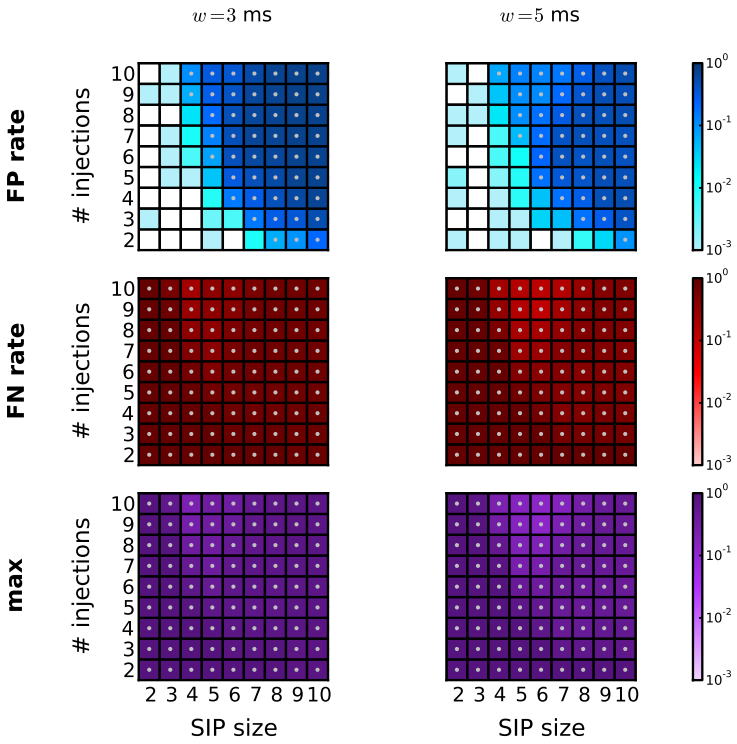

Figure 2.9:
Performance of PSR under jittered synchrony. Performance of PSR (combined filtering with $h = 1$, $k = 2$) on data from SIP models with jittered synchrony. The spikes of SIP events are randomly jittered randomly up to ± 1 ms around the original occurrence time. The performance is shown in terms of FP rates (first row), FN rates (middle row) and maximum of the two (bottom row) for different bin widths: $w = 3$ ms (*left column*) and $w = 5$ ms *(right column)*.

2.3.4 Jittered synchrony

So far we have implicitely assumed that synchrony occurs at the time resolution of the data, i.e. that spike times of the assemblies are not jittered in time. In electrophysiological data this is a rare scenario, and instead spike synchrony typically occurs with a temporal jitter of up to several milliseconds (Grün et al., 1999b; Pazienti et al., 2008). In order to capture such imprecise synchrony, exclusive binning is typically applied (Grün et al., 1999b), where the bin width is chosen large enough to capture the jittered spike pattern. However, the spikes of the pattern may be split into adjacent bins with a probability that depends on the jitter, bin size, and pattern size. Therefore, the original synchronous events are destroyed, leading to increased FN rates (Grün et al., 1999b). Figure 2.9 illustrates how this effect can have a substantial impact on the performance of the method. We applied PSF followed by PSR (combined filtering) on data where synchronous patterns are injected with a jitter $j = \pm 1$ ms, and analysed with a bin width of $w = 3$ ms (left column) and $w = 5$ ms (right column). The performance drops considerably due to an increase of the FP rate for higher z and c, and an overall increase of the FN rate. The performance is slightly better for a bin width of 5 ms, but still unsatisfactory. Consistent with these findings, Grün et al. (1999b) showed that for two parallel spike trains about 60% of the synchronous events are lost if the bin width corresponds to the jitter width. An earlier modification of exclusive time binning (multiple shift method, Grün et al., 1999b) that avoids the splitting of jittered synchrony is not trivially applicable to large numbers of parallel spike trains.

Borgelt and Picado-Muiño (2014) recently implemented a method for pattern detection, called COCONAD (COntinuous time ClOsed Neuronal Assembly Detection), based on the inter-spike distances rather than discrete time binning. The method can be embedded in the FIM methodology and classifies a group of spikes as synchronous if those spikes

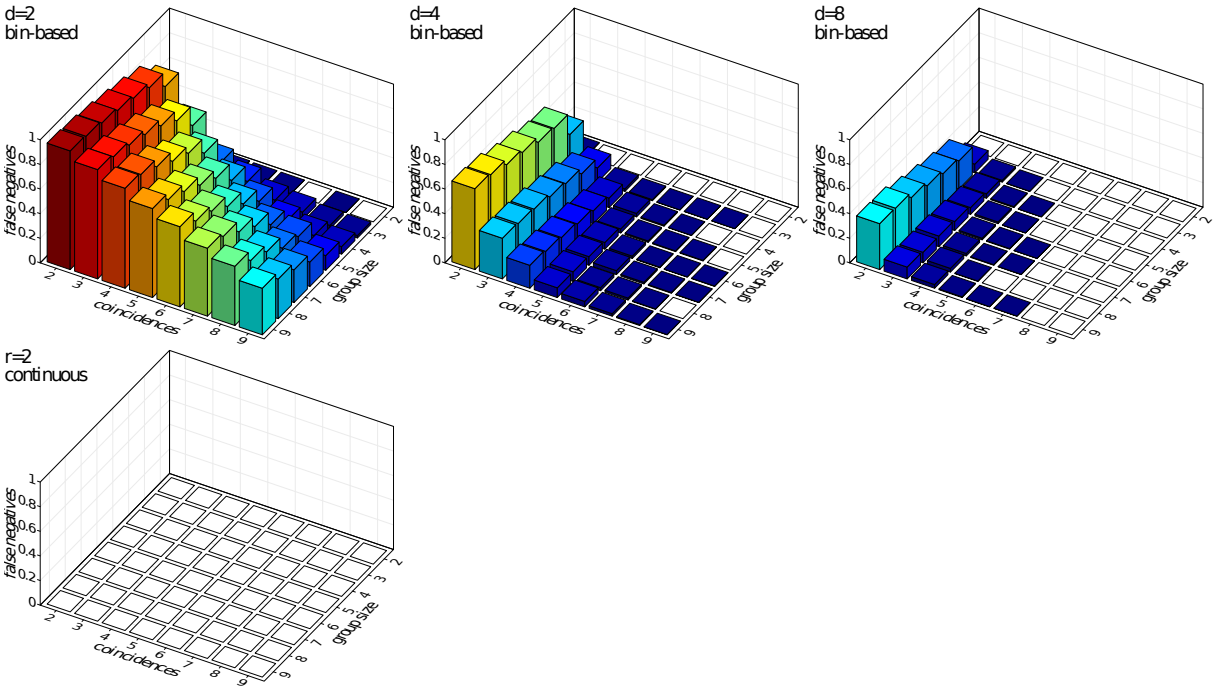

Figure 2.10: **Performance of time binning and COCONAD in the presence of jittered synchrony.** Each panel shows the fraction of jittered synchronous patterns ($j = \pm 1$ ms) of a certain size (x-axis) and support (y-axis) that are injected in otherwise independent Poisson data but are not identified by binning (*top*) or by COCONAD (*bottom*), and thus yield FNs (z-axis) prior to any statistical assessment. Analysis and figure created by David Picado-Muinõ.

fall within a chosen maximum distance w (counterpart of the bin width) from each other. Thus, it identifies patterns with 100% accuracy as long as w is larger than the maximum jitter in synchrony, as illustrated in Figure 2.10.

COCONAD can be trivially substituted to exclusive binning in counting pattern occurrences. Therefore, in applications to data where jittered synchrony is supposed to be present, as in most real-case scenarios, we suggest to count pattern occurrences both in the original and jittered data using COCONAD rather than time binning.

2.4 Discussion

In this study we have presented SPADE (Spike PAttern Detection and Evaluation), a method to detect significant patterns of synchronous spikes extracted from all possible subsets of spike trains in MPST data, in the presence of background activity. Our work is rooted in Picado-Muiño et al. (2013), who demonstrated how to efficiently detect spike patterns in such data, and how to assess their significance under the null hypothesis of spike train independence by a procedure called pattern spectrum filtering (PSF). Here we refined PSF by turning it into an upper-tail test, which is more suitable in applications to data from large numbers of spike trains. Beside, we noticed that PSF is prone to report FP patterns that arise due to the activation of an actual assembly mixed with chance synchrony because of background activity. To identify and remove these FP detections, we introduced here PSR as an additional statistical testing step. As shown in Figure 2.6 (second to last columns), PSR succeeds in eliminating FPs for a wide range of parameters, at the expense

of a minor increase in FNs. Calibration of various types of test data demonstrates the effectiveness an reliability of SPADE in the presence of typical features of real data such as firing rates which are time-varying and heterogeneous across neurons.

The relevance of higher-order correlations (HOCs) for information processing in the nervous system is hotly debated. Approaches based on maximum entropy models, such as Schneidman et al. (2006), suggest that HOCs contribute by a negligible fraction to the total network correlation, which appears to be dominated by pairwise correlations. However, these studies estimate the overall magnitude of that correlation order, and are not sensitive to individual correlation structures of that order. Thus, the presence in the data of a single group of correlated neurons with a certain size is not enough for maximum entropy models to report significant correlation of the corresponding order. The study by Shlens et al. (2006) addresses this point, discussing that maximum entropy models may miss higher-order correlations because they overall contribute only by a negligible fraction to the total correlation. Besides, Roudi et al. (2009) showed that the statistical power of maximum entropy models describing spike correlations in heavily undersampled biological systems (such as parallel recordings with electrode arrays) is low. In Appendix A we show that even without sub-sampling and in the presence of low-dimensional data, HOCs fully contains in the data may not be captured by the total network correlation as estimated by Schneidman et al. (2006). Despite these challenges, Ohiorhenuan et al. (2010) showed using a maximum entropy model approach that in visual cortex local microcircuits exhibit evidence of higher-order interactions, whereas correlation statistics across long-range connections are explained on the basis of pair-wise interactions. However, methods designed to investigate individual spike patterns are needed to investigate the detailed structure of correlation in groups of spiking neurons.

A majority of current methods for spike correlation analysis limit themselves to fully synchronous patterns or to patterns with a specific size of typically low order (e.g. Grün et al., 2002a,b; Berger et al., 2007, 2010; Shimazaki et al., 2012). Other approaches, such as CUBIC (Staude et al., 2010b), conclude on the presence of higher order correlations based on the statistics of the population activity without identifying the specific units engaged in such correlations. While Gansel and Singer (2012) presented a method for the detection of higher-order patterns, they identify pattern subsets by a purely heuristic procedure that is not accessible by analytic treatment, and that tests patterns directly, thus requiring a strong statistical corrections to avoid FPs in MPST (at the expense of FNs). SPADE instead first tests the significance of pattern signatures. PSF eliminates non-significant signatures based on surrogate data through the significance spectrum (see Figure 2.4), and determines the class $\mathcal{P}$ of associated significant patterns. Testing patterns on the basis of their signature rather than testing individual patterns reduces the number of required statistical tests to the number of signatures found in the data. We have shown that the composition of assembly and background spikes typically leads to the identification of additional significant patterns (i.e., FPs). In order to remove this type of FPs, we introduced here the PSR procedure that is based on conditional pairwise tests.

We tested the performance of SPADE on artificial data containing multiple occurrences of a pattern of synchronous spikes injected in the background activity of otherwise independent Poisson spike trains (SIP model, cf. Kuhn et al. (2003a)). We studied the rate of FP and FN detections for occurrence rates of the synchronous pattern varying from 0.66 to 3.33 Hz, which reflect plausible values for the activation frequency of the assumed assemblies (Grün et al., 1999b; Denker et al., 2010). The calibration shows in particular that by introducing PSR, assembly detection becomes possible with near perfect reliability and precision for a large range of SIP parameters. Statistical significance requires higher support and assembly size as the bin width or the neuronal firing rates increase, as illus-

trated in Figure 2.4. Nevertheless, for physiologically realistic parameters only for very small or very infrequent SIP injections these patterns cannot be distinguished from chance synchrony. Moreover, evaluating patterns obtained from a larger set of simultaneously recorded neurons will have only minor impact on our findings due to a slight increase in the average size of observed patterns.

Non-stationarities of the firing rate in time or across neurons are a common concern faced by correlation analysis methods. The effect of non-stationary firing rates on PSF is two-fold. First, the surrogates used to calculate the significance estimates on pattern signatures should adequately reproduce the experimental rate profiles. Even if the underlying rate profile is not known, a variety of suitable approaches for surrogate generation is available for this task (Grün, 2009; Louis et al., 2010). However, the sensitivity in the detection of assembly activations further depends on where do these occur with respect to the rate non-stationarity. In this respect we tested the performance of PSF and PSR in a scenario of step-wise non-stationary firing rates where spike patterns were injected at selected rate levels only. Compared to the stationary case, the method retains a high performance for large parameter regimes (Figure 2.8), and shows only a slight increase in the number of FNs. For very large rate non-stationarities, a time-resolved analysis may be used to additionally aid the detection, as done e.g. in the Unitary Events analysis (Grün et al., 2002b). In a similar framework, we found that also heterogeneous firing rates across neurons (Figure 2.7) exhibit a performance similar to the stationary case. While we see minor increases in the number of FPs, we highlight that to a large extent these are indeed supersets of the injected pattern due to the high probability of gaining an additional coincident spike by chance from the set of neurons spiking at high rates.

For simplicity, we assumed in most of this chapter that synchrony between spikes is precise, i.e. occurs with no jitter, and we identified synchronus events by time binning. This is a non-realistic scenario in most applications to electrophysiological data, where - even assuming a temporal code - spike times are to some extent noisy. Time binning often fails to detect synchronous events even when their jitter is smaller than the bin width (see Figure 2.9). This is particularly true for synchrony among many neurons. However, time binning can be trivially replaced with a time-continuous pattern identification called COCONAD (Borgelt and Picado-Muiño, 2014), which classifies groups of synchronous spike based on the maximum distance between all spikes in the group. By doing so, the method is able to identified jittered synchrony which time binning would likely miss (Figure 2.10). COCONAD uses the technology offered by FIM to implement an effective search for frequent, closed patterns, and can therefore effectively replace time binning in the analysis pipeline of SPADE.

A further scenario that remains to be addressed in the future is the unreliability of spiking activity that causes neurons to selectively skip participation in assembly activations. This scenario was discussed in the context of the synfire chain model, where it was shown that stable propagation of synchronous spike packages through the network happens reliably although the probability that individual neurons participate in each activation of the synfire chain is lower than 1 (Diesmann et al., 1999). Selective participation may arise as a consequence of synaptic failure. The multiple interaction process (MIP; Kuhn et al., 2003a) was proposed as a stochastic model implementing such a behavior. Our method would interpret the variable composition of spikes in a single MIP event as occurrences of multiple SIP events of lower support.

We conclude with a discussion of the practical application of SPADE on data from electrophysiological recordings. Given a set of parallel spike recordings obtained at a resolution (i.e., binning) w, we choose the minimum pattern size z_0 and the minimum pattern support c_0 of the pattern analysis. First, the spike data is binned and, using FIM,

the CFISs and the corresponding pattern signatures are obtained from the transaction list. While this approach is feasible for the experimental data available today, with several hundreds of parallel recordings the computational effort may become too large. In this scenario, we suggest to pre-filter the data entering the analysis as suggested by Berger et al. (2010) before applying FIM on the reduced set of neurons. To monitor dynamic changes in the correlation structure of the activity, e.g., if assemblies are time locked to a particular behavioral event, one may choose to additionally perform the analysis in sliding windows.

Next, the significance of the observed patterns is evaluated by PSF under the null-hypothesis of full independence implemented by uncorrelated surrogate data. For experimental data, several techniques for surrogate generation based on stochastic sampling have been proposed in the past (for a review, see Grün, 2009). Surrogates that preserve the firing rate profiles, such as spike dithering, seem most appropriate since PSF determines pattern significance based on thefiring rates. Given the significance level α and m detected pattern signatures, a minimum of $K = \lceil m/\alpha \rceil$ surrogates are required to achieve the Bonferroni-corrected significance level $\alpha_* = \alpha/m$. Once the surrogates have been generated, we follow the procedure described for the simulated data. CFISs, pattern signatures and the resulting binary pattern spectrum are obtained for each surrogate run. Next, the p-value spectrum is obtained as an average of the binary spectra (see Section 2.1.2). The signatures whose p-values do not exceed the Bonferroni-corrected significance level α_* are marked as significant, and the CFISs of significant signatures are collected into the class $\mathcal{P}$ of potential assemblies. Finally, PSR with combined filtering is performed to reduce $\mathcal{P}$ to a subclass $\mathcal{Q}$ of patterns which are mutually significant with respect to each other.

In summary, SPADE represents a powerful tool to extract candidate cell assemblies from experimental data. The method is statistically rigid, computationally feasible, robust against heterogeneity in the data, and powerful enough to deal with the limited amount of data typically available from electrophysiological experiments.

In Chapter 3 we employ SPADE to identify patterns of synchronous spikes in motor cortex data from two monkeys engaged in an instructed-delay reach-to-grasp task, and to related the presence and features of the found patterns to behavior.

Chapter 3

Synchronous spike patterns in macaque motor cortex during an instructed-delay reach-to-grasp task

In Chapter 2 we presented SPADE, a novel method to analyze patterns of synchonous spikes in MPST data. Here we employ SPADE on data from two macaque monkeys engaged in an instructed-delay reach-to-grasp task to determine the emergence of spike synchronization in relation to behavior. The data are recorded from each monkey via a 100-electrodes Utah grid implanted in motor cortex, over 10 sub-sessions per monkey spanning a period of several months. In sub-session is analyzed individually, separately for each of four trial types and six time windows (*epochs*). We characterize the spike patterns found during each trial type and epoch in terms of neuronal composition, size, occurrence count and spatial arrangement of the composing neuron on the electrode grid. By performing the analysis separately on each of four different trial types and six consecutive time windows, we are able to relate the mentioned pattern features to different behavioral contexts.

This work has been submitted for publication to the Journal of Neuroscience and is currently under review (Torre et al., 2015b). Emiliano Torre designed the analysis of the data under the supervision of Michael Denker and Sonja Grün, and performed the analysis together with Pietro Quaglio. Alexa Riehle and Thomas Brochier provided the data and suggested the analysis of pattern orientation.

3.1 Data and data pre-processing

3.1.1 Experimental protocol and experimental setup

We analyze electrophysiological data recorded from the motor cortex of two monkeys (*macaca mulatta*), monkey L and N, during repeated execution of an instructed-delay reach-to-grasp task. The experimental protocol is described in Riehle et al. (2013). In summary, the monkeys use one of 2 grips - either a full-hand side grip (SG) or a two-fingers precision grip (PG) - to grasp and pull a cubic object along the horizontal axis. The object load is computer-controlled and was either 1 N for low force (LF) or 2 N for high force (HF) trials, respectively. The monkey has to hold the object in a narrow position window for 500 ms to obtain a food reward (drop of apple juice). Task instructions are provided by five LEDs (four red LEDs placed at the corners of a square and a yellow LED in the center) located just above the object. The task is illustrated in Figure 3.1. The monkey has to close a switch located at waist level to self-initiate the trial. After 400 ms

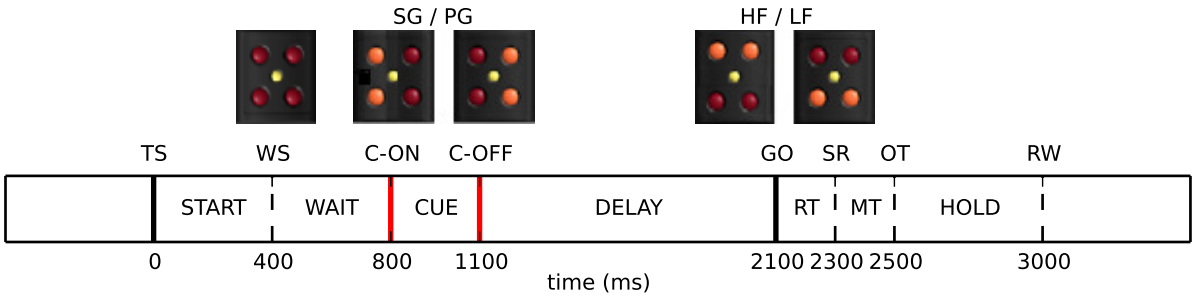

Figure 3.1: **Reach-to-grasp experimental protocol.** The trial start (TS) is self-initiated by the monkey by closing a switch located at waist level. A warning signal (WS) prepares the monkey for the visual cue (C+ until C−) instructing it about the grip type to use: SG or PG. 1 s later a second visual cue turns on (GO), specifying the force needed to pull the object (HF or LF) and requesting movement initiation. The movement onset is marked by the switch release (SR). After object touch (OT) the monkey holds the object for 500 ms in a narrow position window until reward (RW). The timing of the events following the GO signal varies depending on the monkey's reaction time and movement speed. Representative values are indicated in the figure for SR, OT and RW.

from trial start (TS), the yellow LED is illuminated as a warning signal (WS) for the cue occurring 400 ms later. The cue (illumination of the two right or the two left LEDs) is presented for 300 ms and instructs the monkey about the type of grip to be used to grasp the object. After a preparatory delay of 1000 ms, a second cue (illumination of the two top or the two bottom LEDs) indicates the object load and also represents the GO signal instructing the monkey to execute the movement. Therefore, the monkey can perform one of 4 possible trial types: SGHF, SGLF, PGHF, PGLF. The 4 trial types are presented in random succession within a recording sub-session, which consists of a sequence of about 120 consecutive trials executed for about 15 min.

Both monkeys are chronically implanted in motor cortex with a 100-electrodes Utah array (Blackrock Microsystems, Salt Lake City, UT, USA), covering part of the dorsal premotor (PMd) and the primary motor (M1) cortex (see Figure 3.7A for array locations). The length of the electrodes is 1.5 mm, with an inter-electrode distance of 400 μm. Data are recorded using the 128-channel Cerebus acquisition system (Blackrock Microsystems). The signal from each active electrode (96 out of the 100 electrodes were connected) is pre-processed by a head stage with unity gain and then amplified with a gain of 5000. The signal is sampled at 30 kHz and filtered in two different frequency bands to be split into local field potentials (LFP, $0.3 - 250$ Hz) and spiking activity ($0.5 - 7.5$ kHz in monkey N and $0.25 - 7.5$ kHz in monkey L). The potential spike times are identified online on every channel by a threshold crossing algorithm (Blackrock NSP) and the corresponding waveforms saved in the Blackrock Central Suite as snippets of 1.6 ms in monkey L and 1.3 ms in monkey N around the spike time. All behavioral data such as stimuli, switch release, force traces for thumb and index fingers and object displacement are fed into the Cerebus system, sampled at 1 kHz and stored for offline analysis.

3.1.2 Hyper-synchronous events and signal-to-noise ratio

A preliminary analysis of the data reveals for both monkeys and in each sub-session the presence of synchronous spikes at the sampling resolution $\delta = 1/30$ ms of the recording system, a phenomenon invisible to bare visual inspection of the population raster plots. These spikes survive both the automatic noise-cleaning procedure by the Blackrock system

and the offline spike sorting procedure detailed in Riehle et al. (2013). The number of SUAs involved in these hyper-synchronous events ranges from 2 to over 40. Because there are no plausible biological mechanisms of spike synchronization at this high temporal resolution, we investigate whether these hyper-synchronous patterns are expected to occur by chance on the basis of the large number of neurons observed simultaneously and their firing rates. If not, they may be artifacts - possibly due to electrical cross-talks - which the SPADE analysis may classify as false positives.

For each sub-session and for each integer number $\xi \geq 2$, we estimate the expected frequency of hyper-synchronous events of size ξ or higher in the sub-session, under the null hypothesis that the spike trains are independent at the temporal scale $\delta = 1/30\,\mathrm{ms}$ which defines hyper-synchronization. To simplify the math, we consider each spike train as a stationary Poisson process with a firing rate given by its average rate over the full sub-session.

A neuron k, $k = 1, 2, \ldots, N$, that fires $\hat{n}_k$ total spikes in the observation period T has average empirical rate $\hat{\lambda}_k = \hat{n}_k/T$. Under the hypothesis of poissonianity, the probability for that neuron to fire in a given time bin of width δ (the allowed range for hyper-synchronization) is $p_k = 1 - \exp(-\delta \cdot \hat{\lambda}_k)$. The joint probability $p(\xi)$ that exactly ξ out of N neurons fire in that time bin, each with its own probability p_i, is thus a Poisson Bernoulli distribution

$$p(\xi) = \sum_{A \in \mathcal{P}_{N;\xi}} \left\{ \prod_{k \in A} p_k \prod_{h \in A^C} (1 - p_h) \right\}, \tag{3.1}$$

where $\mathcal{P}_{N;\xi}$ is the family of all possible subsets of ξ elements that can be extracted from the set $\{1, \ldots, N\}$, A is one such subset, and $A^C = \{1, \ldots, N\} \backslash A$ is the complement of A. The expression represents the sum, over all possible groups A of ξ out of N neurons, that a specific group A spikes in the bin while the rest of the neurons, A^C, does not.

The sum involves $\binom{N}{\xi}$ terms for each ξ, for a total of 2^N terms. The calculation is computationally infeasible for large N, but can be approximated, in light of Le Cam's theorem, by a Poisson pdf with probability parameter $\sigma = \sum_k p_k$ (see Le Cam, 1963)

$$p(\xi) \approx \tilde{p}(\xi) := \frac{\sigma^\xi e^{-\sigma}}{\xi!}$$

Figure 3.2A shows the difference between the true Poisson Bernouilli distribution $p(\xi)$ and its Le Cam's approximation $\tilde{p}(\xi)$ in two cases, corresponding respectively to probabilities $p_k = 3.33 \cdot 10^{-4}\, \forall k$ and $p_k = 3.33 \cdot 10^{-5} \cdot k$. These two scenarios correspond to spike trains with rates $\hat{\lambda}_k = 10\,\mathrm{Hz}\, \forall k$ and $\hat{\lambda}_k = k\,\mathrm{Hz}$, segmented in time bins of width $\delta = 1/30\,\mathrm{ms}$ to determine hyper-synchrony. In both cases, we determine the true distribution $p(\xi)$ with a Monte-Carlo approach, as the histogram over a sample of 10^8 elements obtained as sums of 100 Bernoulli random variables with probabilities p_k. Both scenarios yield an integrated absolute error < 0.001, demonstrating the validity of Le Cam's approximation to estimate the likelihood of hyper-synchronous events in the presence of typical values of firing rates, also when rates are strongly heterogeneous across neurons.

The cdf $P(\xi) = \sum_{l \geq \xi} \tilde{p}(l)$ provides the approximate probability that a given bin contains an hyper-synchronous event of size ξ or larger. This theoretical probability can be compared to the empirical one, calculated from the data as the fraction of bins of width δ containing hyper-synchronous events of size $\geq \xi$. The comparison is shown in Figure 3.2B.

The analysis run separately of all 10 sub-sessions considered for each monkey reveals that hyper-synchronous events of size 3-4 or larger occur in the data considerably more

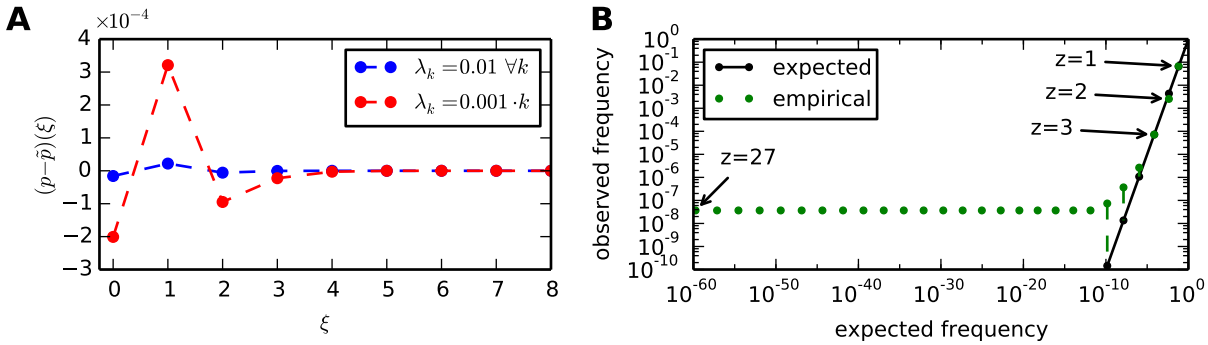

Figure 3.2: **Probability of occurrence of hyper-synchronous events. (A)** Difference between exact and approximate probability mass functions (pmfs) for two Poisson Bernoulli distributions with identical probabilities $p_k = 3.33 \cdot 10^{-4}\ \forall k = 1, \ldots, 100$ (blue line; integrated absolute error $= 4.6 \cdot 10^{-5}$) and with heterogeneous probabilities $p_k = 3.33 \cdot 10^{-5} \cdot k$, $k = 1, \ldots, 100$ (red line; integrated absolute error $= 5.4 \cdot 10^{-4}$). **(B)** Empirical (blue dotted line) vs expected (black dash-dotted line) probability that a time bin of width $\delta = 1/30$ ms (sampling resolution of the recording system) contains an hyper-synchronous event of size $\geq \xi$, as a function of ξ.

frequently than expected under independence, while hyper-synchronous events of size $\xi = 2$ occur within the expected frequency range in most subsessions.

A statistical analysis shows that the frequency of their occurrence largely exceeds the level expected considering the observed population firing rate. Furthermore, fine-scale temporal correlation in the activity of the recorded neurons cannot explain these hyper-synchronous events, as it is not expected at sub-millisecond precision. Assuming that these events, whose presence heavily affects any analysis of spike synchronization, may be artifacts of unknown origin, we adopt a conservative approach and remove all of them from the data we then analyze with SPADE. This may result in the removal of real synchronous spikes, leading thus to false negative (FN) results. However, we privilege this compromise over the risk to have false positive (FP) outcomes. We additionally remove each spike occurring within a time interval $\delta = 1/30$ ms from any hyper-synchronous event, as they may reflect noise too.

For further analysis, we additionally selected only SUAs whose waveform's signal-to-noise ratio (SNR) is > 2.5, as suggested by Hatsopoulos et al. (2004).

3.2 Identification and characterization of synchronous events

3.2.1 Identification of synchronous spike patterns by SPADE

We employ the SPADE analysis presented in Chapter 2 to examine our data for patterns of synchronous spikes occurring more often than expected by chance under the null hypothesis H_0 of independence. MPST data typically contain large numbers of potential patterns to be found and evaluated for significance. SPADE overcomes the severe computational resulting from having to count the occurrences of each such pattern, and the multiple testing problem of assessing their statistical significance.

We define spike synchrony using COCONAD, the time-continuous version of frequent item set mining that does not rely on time binning to find ensembles of synchronous spikes

(see Section 2.3.4). All groups of spikes within a maximum temporal distance w from each other are classified as synchronous. We set $w = 3\,\text{ms}$.

We perform pattern spectrum filtering (Section 2.1.2), which evaluates the p-value of the observed pattern signatures $\langle z, c \rangle$ (where z is the pattern size and c the number of pattern occurrences) over $K = 10,000$ surrogates. Surrogate data are obtained from the original data by dithering each spike time around its position by a random amount, up to $\pm 25\,\text{ms}$. Signature significance is assessed at a level $\alpha = 0.01$, corrected by FDR correction. Finally, we perform pattern set reduction (Section 2.2) to remove conditionally non-significant patterns. Specifically, we opt for combined filtering (Section 2.2.2) with parameters $h = 1$, $k = 2$.

The result of this procedure is a set of patterns which i) are frequent, i.e. occur at least twice, ii) are closed, i.e. repeat more often than their supersets and are in this sense non-trivial, iii) are composed of a larger number of neurons and repeat more times than one may expect by chance in the data under consideration of the empirical firing rates and under the assumption of independence, and iv) are conditionally significant given each other.

3.2.2 Pattern specificity to behavior

Our data consist of parallel spike trains recorded during four different trial types (resulting from the combination of two grip modalities and two force levels). To associate the emergence of significant patterns to behavior, we additionally split each trial into six temporal epochs of 500 ms each (see Section 3.3) and search for significant patterns separately in each combination of epochs and trial types. This results in 24 data sets per sub-session (one per epoch and trial type). We then define four behavioral contexts: trial epoch, trial type, grip modality and force level, each having multiple instances (six epochs, four trial types, two grip modalities and two force levels, respectively). For each behavioral context (e.g. epoch) and each of its instances (e.g. the *start* epoch) we determine whether the significant patterns found during that instance occur exclusively during that instance, or whether instead they occur during other instances of the same behavioral context (e.g. during other epochs).

Formally, given a context with n instances let $\mathcal{P}_i$ be the ensemble of significant patterns found by SPADE during the i-th instance, $i = 1, \ldots, n$. (For example, let P_i, $i = 1, 2, \ldots, 6$ be the ensembles of patterns found during each of the 6 epochs). For each instance containing at least one significant pattern (i.e. $\mathcal{P}_i \neq \emptyset$) we define its pattern specificity as the fraction of patterns in $\mathcal{P}_i$ that do not belong to any other $\mathcal{P}_j \neq \mathcal{P}_i$,

$$S_i = \frac{|\mathcal{P}_i \backslash (\cup_{j \neq i} \mathcal{P}_j)|}{|\mathcal{P}_i|}, \tag{3.2}$$

(for example, the fraction of patterns found during the *start* epoch that do not occur in other epochs). S_i takes values in the interval [0, 1]. It takes value 1 (or: 0) if all patterns found in the associated instance do not occur in any (or: occur in at least one) other instance of the same context.

3.2.3 Pattern clustering

Clustering n patterns into k clusters. Patterns found as significant by the SPADE analysis may overlap in their neuronal composition. We characterize their similarity by performing a clustering analysis based on their neuronal composition. We first define the distance $d(A, B)$ between any two patterns A and B as their Ochiai distance (Ochiai, 1957)

$$d(A,\, B) = 1 - |A \cap B| / \sqrt{|A| \cdot |B|},$$

which takes values between 0 (identical sets) and 1 (disjoint sets). We then employ the matrix D of distances $d(A,\, B)$ between any two patterns as the metric in a classical k-medoids clustering algorithm (Kaufman and Rousseeuw, 1987). By fixing the number k of clusters into which n patterns have to be grouped, the algorithm selects k points (here, patterns) as possible cluster centers (*medoids*) and builds clusters C_j, $j = 1, \ldots, k$, by assigning patterns to their closest medoid. The algorithm then computes the distance d_i^* of each pattern from its closest medoid according to the metric D. The overall sum $\sum_{i=1}^{n} d_i^*$ gives an indication of the total "dispersion" of patterns around their closest medoid. The algorithm searches for the combination of k medoids yielding the lowest dispersion $\Delta_k = \sum_i d_i^*$ and returns the associated configuration of clusters as the optimal $k-$clusters configuration $\mathcal{C}_k$ (having optimal dispersion Δ_k). Given the discrete nature of the elements to cluster (sets of points), a k-medoids algorithm is more suitable than algorithms computing averages of the elements, such as classical k-means clustering techniques.

Determining the optimal number k of clusters. Δ_k is a non-increasing function of k: a configuration with $k+1$ clusters and having a dispersion lower than or equal to Δ_k always exists, as long as $k < n$. A trivial example consists in considering a cluster with at least two patterns in the k-clusters configuration and removing from it one non-medoid pattern into a new isolated cluster. This operation removes an addendum from the dispersion of the first cluster and creates a 1-pattern cluster which has necessarily a dispersion equal to 0, yielding an overall lower dispersion. To quantify the optimal number of clusters for a set of n patterns we thus penalized the total dispersion of a configuration by the number of its clusters, defining its *cost* by $R_k = \Delta_k + \rho \cdot k$, $k = 1, \ldots, n$. We then determined the optimal number of clusters as $k^* = \arg\min_k \{R_k\}$, and grouped the patterns into the configuration $\mathcal{C}_{k^*}$. A configuration $\mathcal{C}_{k+1}$ has a lower cost than $\mathcal{C}_k$ - and is therefore to be preferred - only if $R_{k+1} \leq R_k$ (or equivalently, if$\Delta_{k+1} \geq \Delta_k + \rho$), i.e. if adding a cluster makes the cost drop by an amount which is not smaller than the penalty coefficient ρ. We set $\rho = 1$, which also represents the maximum distance between two patterns. This choice ensures that any two patterns in the same cluster share at least one neuron id. Indeed, whenever a cluster C_j in the configuration $\mathcal{C}_k$ contains a pattern A disjoint from all other patterns in C_j (and thus contributing a $+1$ to the dispersion of C_j), there will exist at least one configuration of $k+1$ clusters having a lower cost than $\mathcal{C}_k$: the one where A is an individual cluster separated from C_j.

3.3 Results

3.3.1 Pattern significance across epochs and trial types

We analyze spiking data from multiple single neurons recorded simultaneously during an instructed-delay reach-to-grasp task (see Section 3.1.1) with SPADE (see Section 3.2.1) to detect statistically significant patterns of synchronous spikes. The patterns are significant if they repeat more frequently than expected given the firing rates of individual neurons under the null hypothesis of independence. For each monkey we select 10 sub-sessions which contain trials of four trial types: SGHF, SGLF, PGHF, PGLF (see Section 3.1.1 for more details). Within each behavioral trial we additionally identify six time windows, or epochs, of 500 ms duration each, related to specific triggers and behavioral demands. Each epoch starts at time t_{pre} before a specific trigger and ends at time t_{post} after the trigger, as summarized in Table 3.1 and illustrated in Figure 3.3. Epoch *start* is centered around the

epoch name	trigger	t_{pre}	t_{post}
start	WS	250 ms	250 ms
cue	C+	250 ms	250 ms
early delay	C-	0 ms	500 ms
late delay	GO	500 ms	0 ms
movement	SR	200 ms	300 ms
hold	RW	500 ms	0 ms

Table 3.1: **Definition of trial epochs**. The table summarizes the six different epochs defined for the analysis. Each epoch is a 500 ms time window starting a time t_{pre} before a trigger and ending a time t_{post} after that trigger ($t_{pre} + t_{post} = 500$ ms).

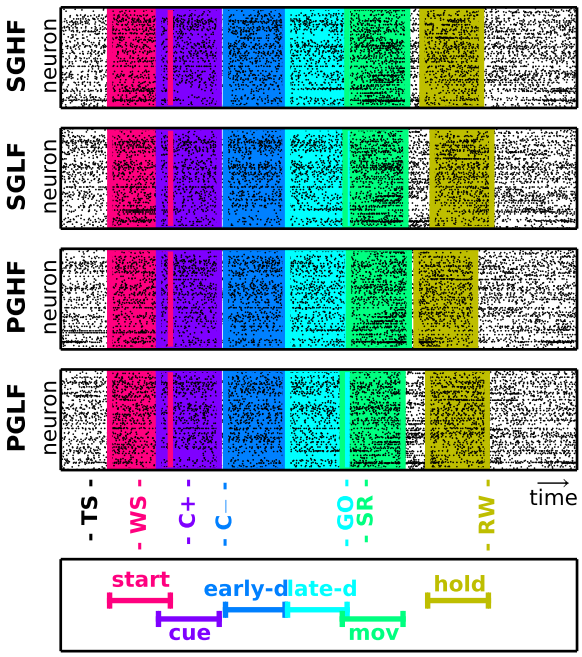

Figure 3.3: **Trial types and epochs.** Each panel shows the simultaneous spiking activity of all neurons (vertical axis) over time (horizontal axis), for trials of the four different trial types from a representative sub-session in monkey N. Each dot indicates one spike. The six colored windows represent the position of the six epochs in the trials. The trigger associated to each epoch and the corresponding epoch name are shown at the bottom (see Table 3.1 for details).

appearance of the warning signal (WS), during a waiting period before the presentation of the visual cue. Epoch *cue* covers the time around the appearance of the visual cue providing prior information about the grip type. Epoch *early delay* captures the first half of the preparatory period after the offset of the visual cue, while epoch *late delay* covers the second half of this preparatory period, until the presentation of the GO signal. Epoch *movement* is centered around the movement onset indicated by the switch release. Finally, epoch *hold* covers the 500 ms time period when the monkey has to hold the object in a fixed position in order to get the reward. Each epoch thus reflects a moment during the trial when different cognitive processes are supposed to happen: paying attention while waiting for the cue (*start*), receiving and/or processing instructions (*cue*, *early delay*, *movement*), preparing the correct grip and waiting for further instruction about the force level (*early delay, late delay*), performing the correct grasp (*movement*) and keeping the position of the object until reward (*hold*). Figure 3.3 illustrates these epochs as colored windows during one selected trial for each of the four trials types, along with the spike times recorded in parallel from all SUAs during the respective trials (black dots). As apparent from the figure, some epochs partially overlap (e.g. *start* and *cue*, or *movement* and *hold*). For epochs following the GO signal, this overlap varies from trial to trial with respect to the reaction and movement time of the animal.

With the goal of relating synchronous spike patterns to behavior, we separate each sub-session into 24 different data sets corresponding to each combination of trial type and epoch. Specifically, each such data set is obtained by concatenating all segments of spike trains corresponding to the same epoch (e.g. start) across all trials of the same type (e.g. SGHF), separately for each neuron. On average, a single sub-session lasts 15 minutes and contains about 130 trials, i.e. over 30 trials of each type. For each sub-session, the individual data sets (e.g. SGHF-start) are concatenated and then separately analyzed with SPADE for significant patterns of synchronous spikes of any size (i.e. number of neurons involved) and neuronal composition. A single subsession contains on average recordings from 70 ± 13.76 single neurons in monkey L and 142.6 ± 14.6 single neurons in monkey N, resulting in millions of different possible patterns which SPADE reduces to a small number of significant patterns. The outcome of the analysis is a collection of 24 sets of significant patterns per sub-session, one per trial type and epoch.

Figure 3.4 shows the results of a representative sub-session from monkey N. Panel A illustrates the composition of neuron identities (horizontal axis) of each significant pattern in this sub-session. 6 patterns are found by SPADE, marked with labels 0 to 5 (vertical axis). Panel B shows a color map for each pattern, representing the p-values of the pattern signature in each epoch and trial type (logarithmic scale). White color corresponds to the significance threshold $\alpha = 0.01$ for pattern signatures before application of the false discovery rate correction (signatures with a lower p-value may eventually still not be significant due to the statistical correction procedure). Squares with a black dot indicate the trial types and epochs where the pattern is classified as statistically significant in the considered sub-session.

3.3.2 Pattern statistics across sub-sessions

Pattern size and occurrences. We perform the analysis outlined above separately for each sub-session. Figure 3.5A shows the total number of significant patterns found across sub-sessions in different epochs (horizontal axis; colors mark the contribution of each trial type). Most significant patterns occur during movement in both monkeys, but all epochs contain significant patterns. No relevant difference can be observed between trial types. Figure 3.5B-C show the average size and number of occurrences of the significant patterns

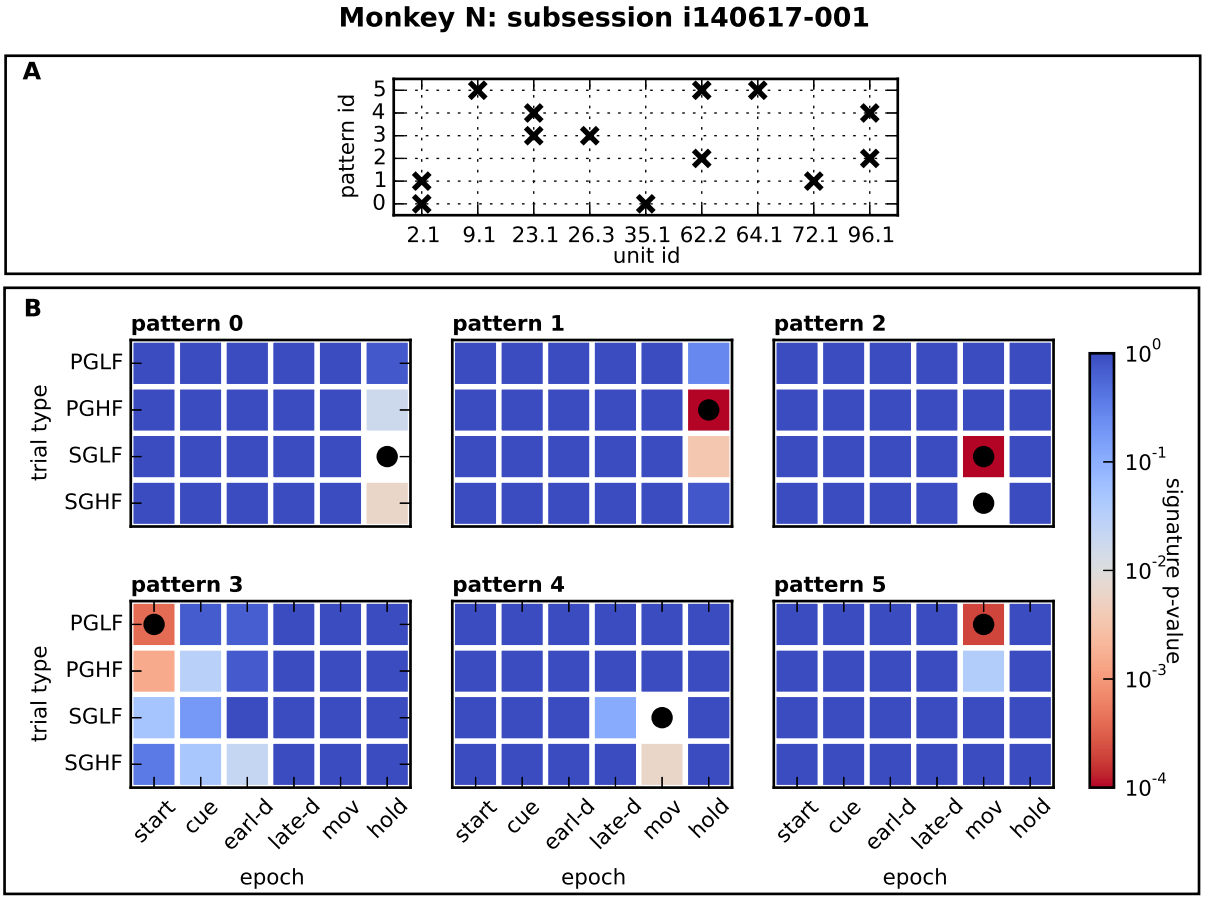

Figure 3.4: **Significant patterns in one representative sub-session. (A)** Composition of neuron identities (horizontal axis) of each significant pattern in the sub-session (vertical axis, pattern ids 0 to 5). **(B)** One color map per significant pattern, showing the non-corrected p-value of the pattern's signature in each trial type and epoch (color progression in logarithmic scale). Squares with the black dot mark the sub-session's trial types and epochs where the pattern is classified as significant.

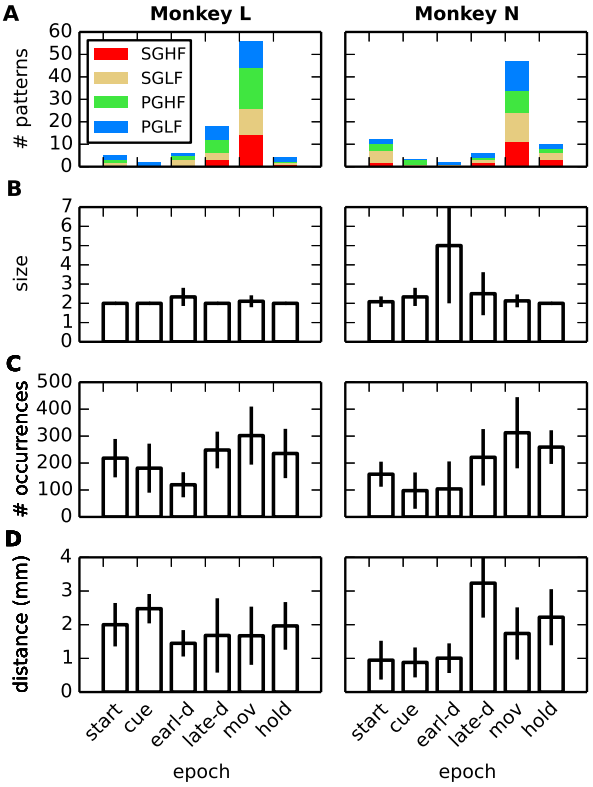

Figure 3.5: **Pattern statistics across sub-sessions.** Various statistics of the significant patterns which are detected across the 10 selected sub-sessions, separately for each monkey. **(A)** The total number of significant patterns in each epoch (colors mark the contribution of each trial type). **(B)** Their average size. **(C)** Their average number of occurrences. **(D)** The average distance between electrodes on which neurons are recorded being involved in the same pattern. Whiskers in panels B to D indicate the standard deviation, when this is larger than 0.

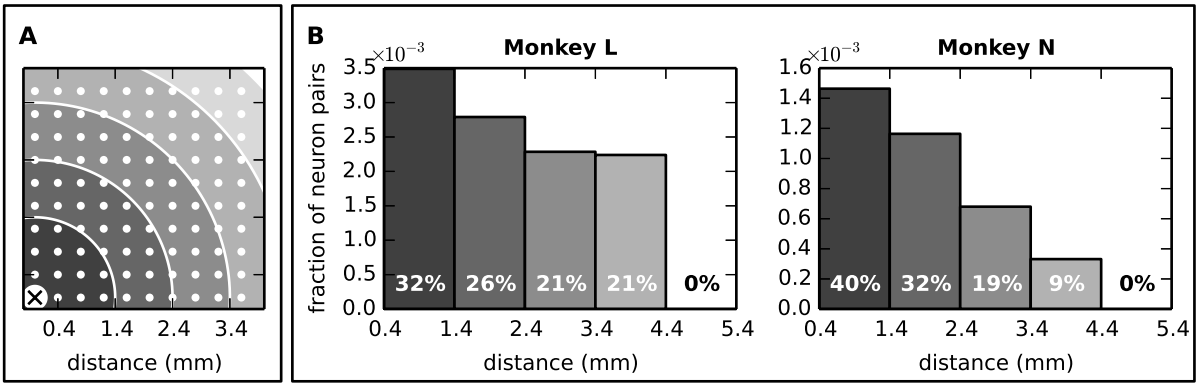

Figure 3.6: **Dependence of synchrony on electrode distance. (A)** Distance of the electrodes on the recording array (white dots) from one reference electrode (cross in the bottom left corner). Dark to light shaded areas cover regions at progressively larger distances from the reference electrode, from 0.4 mm to 5.4 mm in steps of 1 mm. **(B)** Histogram showing, for each distance range, the ratio between the number of neuron pairs involved in the same pattern and placed within that range from each other (summed across patterns found in any sub-session, epoch and trial type), and the total number of recorded neuron pairs placed at that distance (vertical axis). The numbers inside each bar represent the relative height of each bar with respect to the sum of all bars, i.e. the relative fraction of neuron pairs involved in the same patterns placed at a given distance.

detected during each epoch and trial type. Whiskers indicate the standard deviation, revealing a lower variability in pattern size across sub-sessions as compared to the number of pattern occurrences. Most patterns are composed of 2 neurons and repeat 100 to 300 times in each data set. The largest pattern is found in PGLF trials during the early-delay epoch, comprising 8 neurons and occurring twice. The remaining patterns have a size ≤ 3 for monkey L and ≤ 4 for Monkey N. Typically, the number of occurrences needed for a pattern to become significant drops exponentially with pattern size, as shown by Torre et al. (2013).

Spatial arrangement of pattern neurons. The inter-electrode distance on the recording array is 400 μm and 566 μm in the horizontal/vertical and the diagonal direction, respectively. The vast majority of neurons forming significant patterns are recorded from electrodes located 1 to 2 mm apart (Figure 3.5D), corresponding to a separation of 2 to 5 electrodes on the array.

A number of experimental studies showed a decay of spike correlation among neurons placed at increasing distances (Kwan et al., 1987; Gray et al., 1989; Murthy and Fetz, 1996; Berger et al., 2007; Smith and Fetz, 2009). SPADE makes it possible to test this finding on MPST data. Here, we study how the fraction of neuron pairs belonging to any significant pattern varies with the distance between the two neurons in the pair. To this end, i) we extract from each pattern (of any size) found as significant in any sub-session, epoch and trial type all pairs of neurons composing the pattern, and calculate the distance between the neurons in each pair, ii) we build the histogram of the number of pairs placed at a given distance (from 0.4 mm to 5.4 mm in steps of 1 mm), and iii) we correct the values taken by each bar in the histogram by the overall number of recorded neuron pairs placed at that distance, which varies with distance. The results are shown in Figure 3.6. For both monkeys the fraction of pairs of neurons involved in the same pattern

and placed within a given distance range, corrected by the total number of recorded neuron pairs placed at that distance, decays with distance, in agreement with previous findings. The trend is stronger for monkey L than for monkey N, but is present in both animals. Summing across sub-sessions, trial types and epochs, the data from monkey L contain 71 pairs of SUAs involved in the same pattern out of 24905 recorded pairs ($\approx 2.9‰$), while the data from monkey N contain 101 out of $101,204$ correlated pairs ($\approx 0.95‰$). Note that these seemingly low numbers are still considerably higher than the effective significance threshold used in the analysis (1%, corrected for false discovery rate over the number of pattern signatures found in the data - several hundreds - and further decreased by pattern set reduction, which removes chance overlapping patterns. See Section 3.2.1 for details). The overall statistical analysis is very conservative compared to similar studies performed on few simultaneous recordings, which is necessary to avoid FPs.

In order to relate the spatial organization of synchrony to behavior, we additionally investigate whether the location on the array of neurons involved in synchrony changes with the trial type or the trial epoch. Figure 3.7A shows, for monkey L and N, the location of the array on motor cortex with respect to central sulcus (CS), arcuate sulcus (AS) and pre-central dimple (PD; compare with Figure 1 in Riehle et al., 2013). Figure 3.7B shows the total number of SUAs recorded by each electrode, summed over all sub-sessions. In total, 700 SUAs are considered in monkey L (70 ± 13.76 SUAs per subsession) and 1426 (142.6 ± 14.6) in monkey N. Figure 3.7C shows the number of SUAs recorded by each electrode and involved in a significant pattern, summed over sub-sessions and normalized by the number of SUAs recorded by that electrode (shown in panel B). No qualitative differences are found between trial types, whereas for epochs the majority of patterns are found during movement. Neurons involved in significant patterns do not exhibit trial type- or epoch- specific spatial clustering.

Finally, we investigate the spatial orientation of pairs of neurons involved in synchronous spiking activity. For each epoch we extract all pairs of neurons involved in any significant pattern found during that epoch and during any trial type and sub-session, and compute the orientation of the pairs on the grid as angles between $-\pi/2$ and $+\pi/2$. We then build an orientation histogram in polar coordinates, by pooling the angles found in 8 ranges of amplitude $\pi/8$ each. Each bar in the histogram is finally normalized by the total number of recorded neuron pairs with orientation in that range, summed over all sub-sessions and shown in Figure 3.7D. The histogram after normalization is shown in Figure 3.7E for the movement epoch. Qualitatively similar results are obtained for the other epochs. Neuron pairs involved in significant spike patterns show for both monkeys a privileged medio-lateral orientation.

3.3.3 Specificity of patterns to behavioral context

In order to assess whether patterns are specific to behavior, we consider four behavioral contexts: epoch, trial type, grip modality and force level. Each behavioral context has multiple instances: 6 trial epochs, 4 trial types, 2 grip modalities and 2 force levels, respectively. For each such instance i, we quantify its pattern specificity S_i by the fraction of patterns found in that instance (e.g. start epoch) and not in other instances of the same context (e.g. other epochs). $S_i = 1$ and $S_i = 0$ imply maximum and no specificity, respectively.

Starting with the six different epochs from *start* to *hold*, we compute separately for each sub-session the ensembles $\mathcal{P}_i$ of significant patterns found in that sub-session during each epoch i, pooling across all trials irrespective of the trial type. For instance, Figure 3.4B shows that in our representative sub-session the hold epoch contains 2 patterns, named

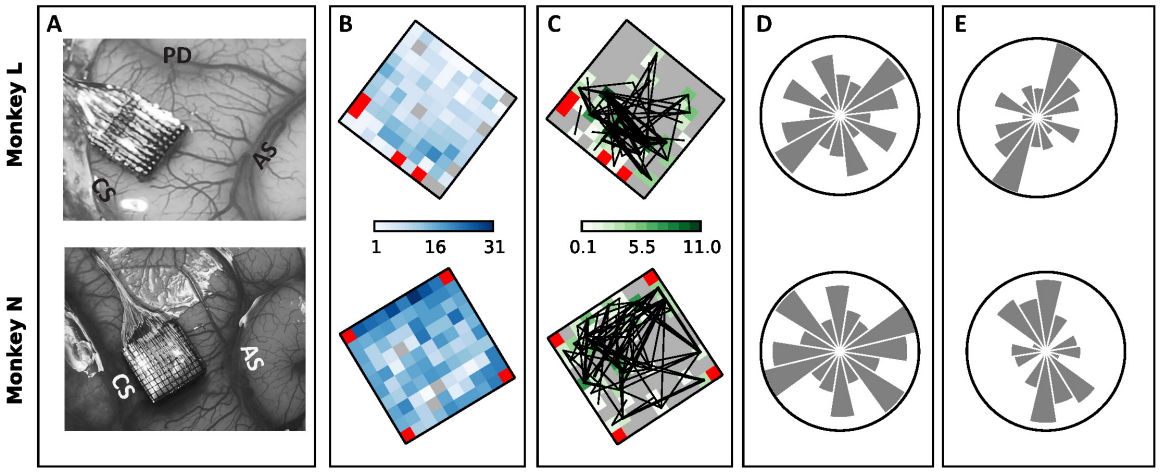

Figure 3.7: **Spatial arrangement of neurons participating in significant patterns.** **(A)** Locations of the electrode arrays in the motor cortex in the right emisphere of the two monkeys. The connector is visible on the upper-left edge. CS: central sulcus, AS: arcuate sulcus, PD: pre-central dimple. **(B)** Color maps of the total number of SUAs recorded on each electrode of the recording array, summed across all selected sub-sessions. Red squares indicate the four unconnected electrodes. **(C)** Color maps showing, for each electrode, the total number of patterns found during the movement epoch (and summed across sub-sessions and trial types) that involve neurons recorded on that electrode, divided by the total number of single neurons recorded by the electrode (panel B). Gray squares correspond to a value of 0, and indicate electrodes that never recorded single neurons involved in significant patterns. Black lines connect electrode pairs extracted from each significant pattern found. **(D)** Spatial orientation of recorded neuron pairs. Each panel shows a bar chart in polar coordinates. The direction of each bar corresponds to one orientation on the recording array, whereas the length of the bar represents the total number of neuron pairs placed along that direction, computed for each sub-session and summed across sub-sessions. The circles have diameter $5,000$ for monkey L and $18,000$ for monkey N. **(E)** Spatial orientation of all pairs of neurons involved in significant patterns found during movement, as a fraction over the total number of recorded neuron pairs with that orientation shown in panel D.

with ids 0 and 1, and that these patterns are only significant in this epoch, which yields pattern specificity 1 for the hold epoch. Similarly, the start epoch and the movement epoch contain 1 and 3 patterns respectively, occurring in those epochs only and yielding again specificity 1. The other epochs of this session contain no significant patterns and a pattern specificity value cannot be assigned to them. For each epoch we then average its pattern specificity across sub-sessions. The results, illustrated in Figure 3.8A separately for each monkey, indicate that most epochs exhibit a high ($> .6$) or very high ($> .8$) pattern specificity, i.e. contain a majority of patterns that do not occur in other epochs. An exception is represented by the cue epoch in both monkeys, and by the early-delay epoch in monkey N, which show specificity 0.5. However, all three cases result from an average computed across 2 sub-sessions only (the other sub-sessions do not contain significant patterns in that epoch) and seem explained by a more general paucity of synchronous events in those time periods. In contrast, the late-delay, the movement and the reward epochs exhibit pattern specificity 1 or close to 1 in both monkeys. Interpreting significant patterns as the signature of cognitive processing which involves the activation of specific cell assemblies, these results suggest that assemblies become selectively active with respect to specific behavioral demands.

We analogously compute pattern specificity for each of the 4 trial types, irrespective of the epoch where they are found. In our example from Figure 3.4B, the pattern specificity for the different trial types in the representative sub-session is 0 for SGHF, 0.5 for SGLF, 1 for PGHF and PGLF. The pattern specificity for the different trial types averaged across sub-sessions, shown in Figure 3.8B, ranges between 0.15 (low) and 0.6 (high). Because a trial type is defined after a combination of the grip modality and the force level required to accomplish the task, we disentangle the two aspects by further computing the pattern specificity for the grip modality (SG vs PG) irrespective of the force level and the epoch (Figure 3.8C), and for the force level (HF vs LF) irrespective of the grip modality and of the epoch (Figure 3.8D). The results show high or very high grip-type specificity (> 0.8) and low force-type specificity (< 0.5). The combination of both explains the moderate specificity values found for the trial types. Following the same line of argument as before, we conclude that in our data specific grip modalities, but not specific force levels, are associated with the activation of specific cell assemblies.

3.3.4 Specificity of pattern clusters

Individual neurons often participate in multiple significant patterns (see e.g. Figure 3.4A). When interpreting patterns as signatures of cell assemblies associated to a specific behavior or condition, different but overlapping patterns may reflect either a common cell assembly, whose activation recruits each time only a subset of its composing neurons, or different cell assemblies which share some of the composing neurons. To distinguish between these two hypotheses, we investigate whether different but overlapping patterns are or are not found during the same behavioral context or not, which support the first and the second hypothesis, respectively. To this end, for each subsession we first group by means of a clustering algorithm all patterns found during that sub-session (across all trial types and epochs) into ensembles characterized by a similar neuronal composition (for details, see Section 3.2.3). Given n total patterns in a sub-session, the clustering algorithm evaluates for each integer $k = 1, 2, \ldots, n$ how well these patterns can be separated into a configuration of k clusters, by assigning a cost to that configuration as illustrated in Figure 3.9B. The algorithm selects the optimal number k^* of clusters and the associated optimal configuration based on the lowest cost.

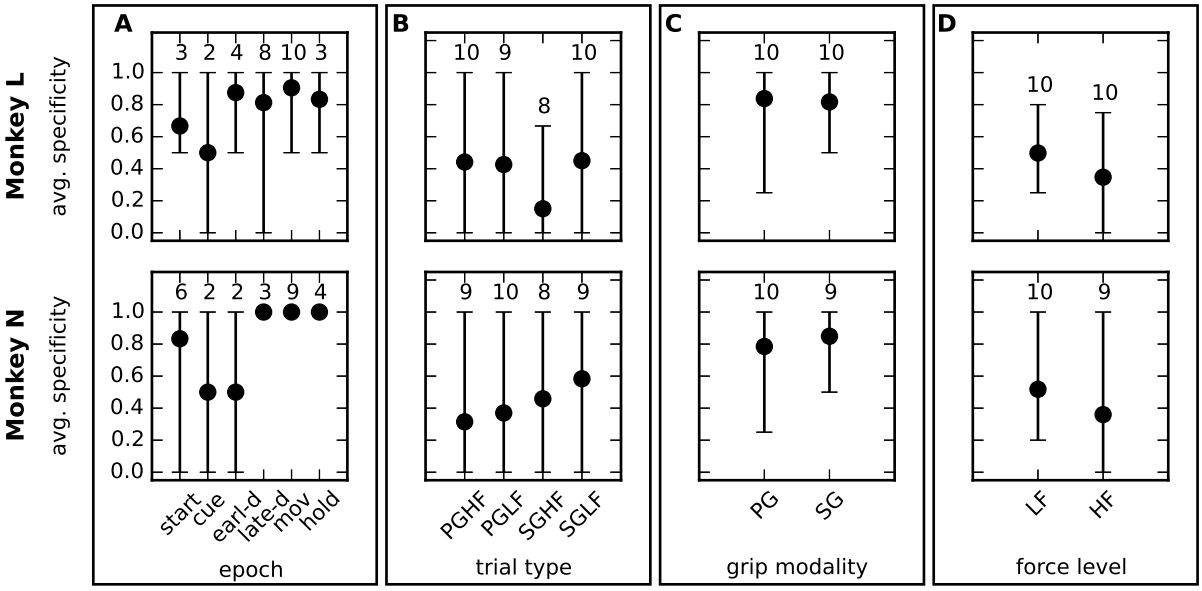

Figure 3.8: **Pattern specificity for different behavioral contexts.** The four panels show the pattern specificity at each instance of the four behavioral contextual conditions: epochs **(A)**, trial types **(B)**, grip modalities **(C)** and force levels **(D)**. The values are calculated for each monkey across all sub-sessions. The bars extend between the minimum and maximum value across sub-sessions. The numbers above each bar indicate the number of sub-sessions which actually contain significant patterns in that instance (e.g. in that epoch). When this number is 2, the values taken by the 2 corresponding subsessions are the ends of the range. The other sub-sessions do not contain significant patterns and a pattern specificity value cannot therefore be calculated.

Figure 3.9A illustrates, for the already discussed representative sub-session of monkey N (compare with Figure 3.4A), how its $n = 6$ patterns are optimally clustered into $k^* = 3$ clusters. Each cluster is identified by a different color. Panel B shows the cost of each configuration of $k = 1, 2, \ldots, 6$ clusters, indeed with a minimum at $k^* = 3$.

Moreover, for each sub-session we compute the cluster specificity of each behavioral context, analogously to what has been done for the pattern specificity. For instance, the cluster specificity of one epoch reflects the fraction of clusters found in that epoch (irrespective of the trial type) and in that epoch only. No value can be assigned to datasets that contain no significant patterns and where no cluster can therefore be identified.

We performed the pattern clustering and computed the cluster specificity of each behavioral context (epoch, trial type, grip modality and force level) separately on each subsession. Figure 3.10 shows the results averaged across sub-sessions for each monkey. Epochs have an overall cluster specificity of 0.3 to 0.5 (panel A), which drops to less than 0.2 for trial types (panel B), grip modalities (panel C) and force levels (panel D). These results indicate that groups of similar patterns are not specific to any of the defined behavioral contexts, contrarily to individual patterns. This supports the second hypothesis outlined above, namely that different cell assemblies as identified by different patterns are associated to different behavioral contexts despite involving partially identical neurons.

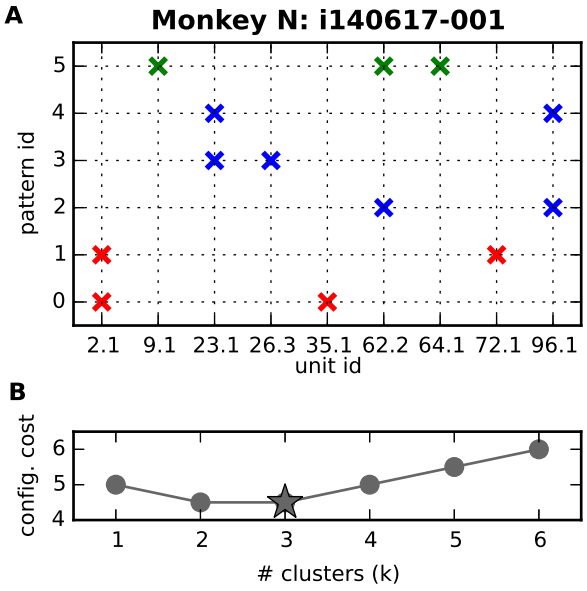

Figure 3.9: **Outcome of clustering for one representative sub-session**. **(A)** The crosses indicate the neuron ids (horizontal axis) of each pattern detected in the sub-session (vertical axis; as in Figure 3.4A). Red, blue and green colors identify the membership of each pattern to one of $k^* = 3$ clusters. **(B)** Cost associated to each number k of clusters, $k = 1, 2, \ldots, 6$.

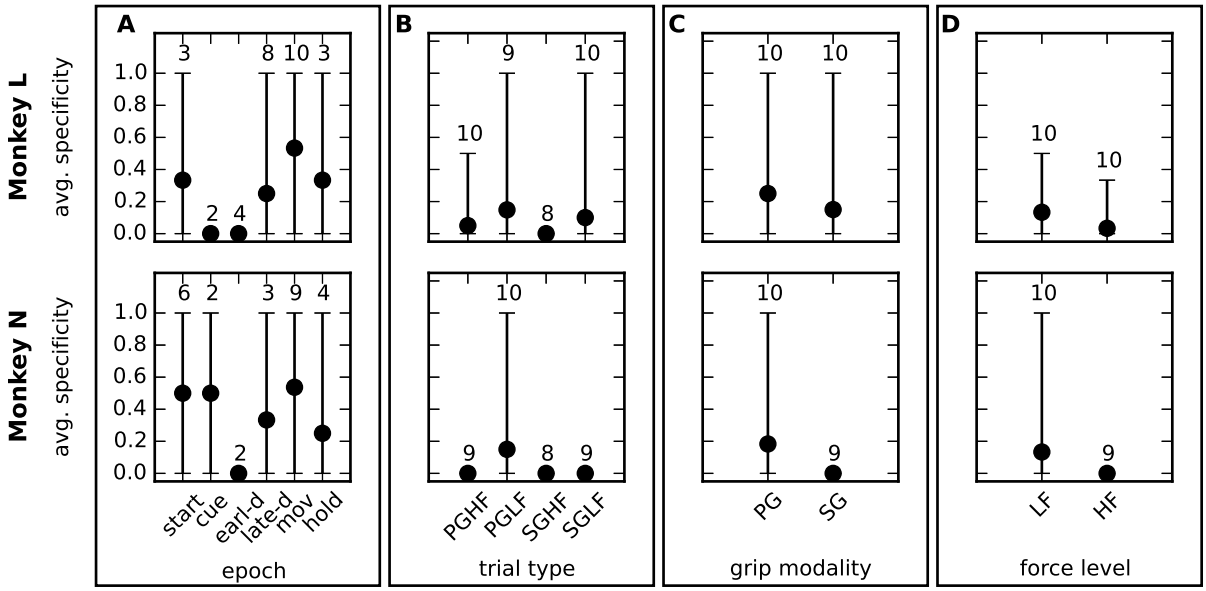

Figure 3.10: **Cluster specificity for different behavioral contexts.** The four panels show the cluster specificity value for epochs **(A)**, trial types **(B)**, grip modalities **(C)** and force levels **(D)**, analogous to Figure 3.8 for pattern specificity.

3.4 Discussion

In this study we examined the emergence of patterns of precise synchronous spiking activity in massively parallel recordings from macaque motor cortex during execution of an instructed-delay reach-to-grasp task (Riehle et al., 2013). To this end we employed the SPADE analysis (Torre et al., 2013), which combines a frequent itemset mining approach to efficiently find repeating patterns of synchronous spikes in large data sets with a hierarchy of statistical tests to determine their significance without incurring into severe multiple testing.

Our analysis revealed the presence of significant patterns in both monkeys. The patterns occurred mainly, but not exclusively, during the movement epoch (Figure 3.5A). The size of these patterns, i.e. the number of composing neurons, was primarily 2, but also 3 or higher (Figure 3.5B). Each significant pattern occurred tens to hundreds of times in each data set (Figure 3.5C). Even though the neurons forming a pattern typically belonged to non-neighboring electrodes (Figure 3.5D), the probability to find 2 neurons belonging to the same pattern decreased with electrode distance (Figure 3.6B). Previous studies also found that synchrony between pairs of neurons decreases with distance (Kwan et al., 1987; Gray et al., 1989; Murthy and Fetz, 1996; Berger et al., 2007; Smith and Fetz, 2009), but with a faster spatial decay, and rarely found synchrony between neurons at a distance larger than 1 mm. However, most of these studies used single electrodes inserted separately into the cortex, which allow for a less precise measurement of inter-electrode distances as compared to the electrode arrays used in this study. Furthermore, these studies were based on the cross-correlation analysis, which accounts for pairwise interactions only. In contrast, the SPADE analysis accounts for the activity of the entire population of neurons under study on the basis of patterns that are expected to be generated by any sub-group of neurons.

We further determined whether patterns are specific to the behavioral context such as epochs, trial types, grip modality and force level. Specificity was defined such that it takes a large value for a behavioral context when a majority of the patterns are detected only in that context (e.g. during the start epoch, or during side grip trials). High pattern specificity for one behavioral context may indicate that the associated patterns are involved in information processing for that specific context. Our results revealed for both monkeys a high pattern specificity with respect to individual epochs (Figure 3.8A) and grip modalities (Figure 3.8C), but not to the force levels (Figure 3.8D). The high specificity for individual epochs matches earlier studies showing that distinct populations of neurons are involved in movement preparation and execution in instructed-delay motor tasks (Wise et al., 1983; Riehle and Requin, 1989; Confais et al., 2012; for a review see Riehle, 2005). The present observation suggests that synchronous spiking activity may play a critical role for the processing of information in motor cortex, either within these different populations or possibly complementary to them. The contrasting degree of specificity between grip types and force levels is likely to reflect functional properties of the cortical location at which the array was implanted. In a previous study (Milekovic et al., 2015) we demonstrated that grip types could be precisely decoded from neuronal data recorded in monkey L, whereas force level decoding was less precise. It seems therefore plausible that synchronous spiking activity may contribute to the fine tuning of neuronal activity for different types of movements, for instance by engaging an assembly of neurons targeting similar sets of muscles (Jackson et al., 2003).

Some of the patterns found in the data shared the same neurons. Therefore, we investigated whether overlapping patterns may be associated to a common cell assembly which required only a subset of its neurons to fire synchronously each time. Otherwise, the presence of overlapping patterns would suggest that the same neuron participates in dif-

ferent assemblies at different times. To discriminate between these two cases, we grouped similar patterns and computed the specificity of behavioral contexts in terms of pattern clusters rather than of individual patterns. The results, shown in Figure 3.10, indicate a very low cluster specificity for all behavioral contexts, supporting the second hypothesis that overlapping patterns indeed reflect the activation of different assemblies.

We also investigated the spatial organization of significantly correlated pairs of neurons in the forelimb area of monkey motor cortex, and found a preferential alignment of the patterns along the medio-lateral axis. This orientation is perpendicular to the dominant orientation reported for traveling waves of beta oscillations in the local field potential (Rubino et al., 2006), or the orientation of the flow of information between individual neurons as revealed by Granger analysis (Takahashi et al., 2015). It is however consistent with the dominant alignment of synchronized pairs of neurons in the motor cortex along the central sulcus (Kwan et al., 1987). This medio-lateral alignment may subserve the tight functional coupling between proximal and distal representations of the upper limb during the performance of complex reach to grasp movements. However, further investigations will be required to confirm this interpretation. The preliminary data obtained partly in the same animals using single intracortical microstimulation (sICMS) on single electrodes on the array while recording from all other electrodes have shown that the preferred orientation between stimulus and recording was primarily parallel to the central sulcus (Hao et al., 2015).

Former analyses of parallel spike trains for excess spike synchrony using the Unitary Events analysis revealed behavior related spike synchrony between pairs or triplets of neurons (Riehle et al., 1997; Kilavik et al., 2009; Ito et al., 2011). Part of these findings were also reproduced by other methods (e.g. Pipa et al., 2008), while Shimazaki et al. (2012) were able to show that spike coincidences between neuron triplets expressed indeed behaviorally locked higher-order correlations. However, these studies were carried on simultaneous recordings from a small number of neurons only. One may expect to observe synchrony between larger groups of neurons as the size of the observed population increases, because the undersampling is reduced. However, reasoning based on combinatorial calculus shows that sub-sampling remains an obstacle to the observation of higher-order correlations. To simplify the calculation, we can assume that the cortical volume in which the arrays we used were inserted contained about $M = 1,000,000$ neurons ($100,000$ neurons per square millimiter [REFS]), $N \simeq 100$ of which we randomly sampled. Assuming that a cell assembly has size C and that its neurons are distributed homogeneously in the volume (or that, given our ignorance about any special spatial arrangement, its neurons are equally likely to be located anywhere in the volume), the probability of sampling exactly ξ assembly neurons is

$$p(\xi) = \frac{\binom{C}{\xi}\binom{S-C}{N-\xi}}{\binom{S}{N}},$$

from which we have

$$\begin{aligned} p(\xi+1) &= \frac{(C-\xi)(N-\xi)}{(\xi+1)(S+\xi+1-N-C)}\,p(\xi) \\ &\approx \frac{NC}{(\xi+1)S}\,p(\xi) \quad \text{for } \xi \ll N, C \ll S. \end{aligned}$$

For instance, the probability to sample $\xi + 1 = 3$ out of $N = 100$ neurons belonging to the assembly of size $C = 100$ would be about 300 times smaller than the probability of observing only $\xi = 2$ of those neurons. $C = 100$ is an interesting value because it represents

in the synfire chain model - a network model where synchrony propagates through links of cell assemblies suitably connected in a chain fashion (Abeles, 1991)- the minimum size of individual links that ensures a stable propagation in the presence of noise (Diesmann et al., 1999). Although based on several simplifications, these numbers give an idea of how much more unlikely it is to observe higher-order correlations in sub-populations of a few hundred neurons randomly sampled from full populations several orders of magnitude larger. In light of these considerations, the few highly significant synchronous spike patterns of size 3 or larger (up to 8) that we found in our data provide a strong indication that larger pools of synchronous cell assemblies may underlie cortical processing. Continuous progress in electrophysiological technology promises to provide in the near future data from thousands or even tens of thousands of neurons in parallel (Schwarz et al., 2014), which are better suited to detect concerted activity of higher complexity.

Chapter 4

ASSET: analysis of sequences of synchronous events in MPST data

With the ability to observe the activity from large numbers of neurons simultaneously using modern recording technologies, the chance to identify sub-networks involved in coordinated processing increases. In Chapter 2 we have presented SPADE, a method to detect statistically significant synchronous spike events in MPST data. Several studies have suggested neural networks models, such as the synfire chain model (Abeles, 1991) able to process and transmit information by propagating synchrony across neuronal groups.

Temporal sequences of synchronous events (SSEs) thus constitute an interesting type of coordinated activity to investigate. Schrader et al. (2008) introduced a visual technique to identify SSEs in MPST, based on an intersection matrix that contains in each entry the degree of overlap of active neurons in two corresponding time bins. Repeated SSEs are reflected in the matrix as diagonal structures of high overlap values. The method as such, however, leaves the task of identifying these diagonal structures to visual inspection rather than to a quantitative analysis.

In this chapter we present ASSET (Analysis of Sequences of Synchronous EvenTs), a fully automated method which rigidly determines the statistical significance of diagonal structures in the intersection matrix and the neuronal composition of the associated SSEs.

This work has been submitted for publication to PLOS Computational Biology and is currently under review (see Torre et al., 2015a). Emiliano Torre has developed the statistical method and designed, under the supervision of Sonja Grün and Michael Denker, the stochastic models employed for calibration. Moritz Helias derived together with Emiliano Torre the first step of the method, i.e. the analytical expression of the probability matrix. Under the supervision of Emiliano Torre, Carlos Canova implemented in his master thesis a previous numerical version of ASSET, and further contributed analyzing the performance of the current method for a varying number of trials. George Gerstein provided feedback and usuful hints during the development of this project.

4.1 Introduction

Synchronous input spikes to a receiving neuron are considered most effective in generating an output spike, as predicted by theoretical studies coining the term coincidence detector (Abeles, 1982b). The argument rests on the premise that excitatory post-synaptic potentials (EPSPs) in the cortex are typically small in relation to the firing threshold, so that many EPSPs need to overlap to produce an output spike. Due to leak currents in the neuronal membrane, the firing threshold is reached with fewer spikes when these

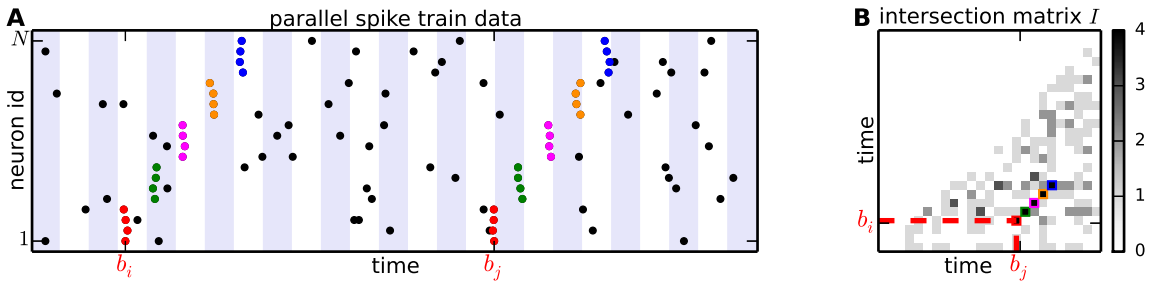

Figure 4.1: **From spike trains to the intersection matrix. (A)** Illustration of parallel spike trains of multiple neurons (vertical axis) over time (horizontal axis). Dots in each row correspond to the spike times of one neuron. Time is discretized into adjacent bins (marked by white and blue shaded backgrounds) to define synchronous events. Synchronous spikes forming anSSE repeating twice are indicated by colored dots (one color per event). **(B)** Intersection matrix I. Each matrix entry I_{ij} (values encoded by gray levels) contains the degree of overlap of neurons active in time bins b_i and b_j. Only the entries I_{ij} with $i < j$ are shown due to the symmetry of the matrix.

arrive synchronously at the post-synaptic neuron rather than sparsely, thus making the neuron behave like a coincidence detector. Experimental studies provide evidence for the existence of coincidence detectors (Roy and Alloway, 2001, see e.g.) and relate them to various mechanisms of spike-timing dependent plasticity (Bender et al., 2006; Fino et al., 2010) as well as to different encoding and decoding schemes (Perez-Orive et al., 2004; Hong et al., 2012).

Cortical anatomy supports such considerations. Individual neurons receive synaptic connections from a large number of neurons (on the order of $10,000$ in the human cortex, see e.g. Braitenberg and Schüz, 1991) and project to a similar number of other cells. Such a connectivity structure combined with suitable synaptic delays may produce spatio-temporal spike patterns, as proposed among others by Abeles (1991), Diesmann et al. (1996), Diesmann et al. (1999) and Izhikevich (2006). A simple case is represented by a temporal sequence of synchronous events (SSE), each event consisting of synchronous spikes from a group of neurons.

The synfire chain (SFC) model (Abeles, 1991) is a neural network model that has been proposed to exhibit such activity through a suitable wiring of the neurons in the network in successive groups interconnected in a highly divergent and convergent manner. Each neuron of one group projects to several neurons of the next group, thus forming a chain structure. Under the assumption that the synaptic transmission delays from neurons of one group to neurons of the next group are identical, the synchronous stimulation of neurons in the first group leads to robust propagation of synchronous spiking activity through the chain (Diesmann et al., 1999) even in the presence of noise. The activation of an SFC would thus lead to the occurrence of an SSE.

In order to assess whether SSEs are indeed observed in the brain and have a functional role, the spiking activity of several neurons needs to be recorded simultaneously, analyzed for temporal correlation and brought in relation to behavior. Under the assumption that an SSE occurs sparsely in a given data set, for instance in relation to a specific behavior, pairwise correlations between neurons engaged in the activity are weak and will not be detected by means of pairwise correlation analyses. For this reason, an analysis method was proposed by Schrader et al. (2008) that directly searches for SSEs in massively parallel spike data.

The basic idea of the method presented by Schrader et al. (2008) is the following: after

discretizing time in bins of a few ms (see Figure 4.1A), a synchronous event which repeats at two different time bins leads to a large number of neurons that are active in both bins. Building an intersection matrix I where each entry I_{ij} represents the number of neurons active in both i and j, a group of synchronous events occurring in bins i and j results in a large overlap I_{ij} compared to other entries of the matrix. An SSE which occurs twice produces a sequence of high-valued entries ,which we name *diagonal structures* (DS), in the matrix I in parallel to the main diagonal (Figure 4.1B). A diagonal filter applied on the matrix I enhances the contrast between the DS and surrounding entries, mapping I into a filtered matrix F.

The method was calibrated using data from large-scale SFC network simulations. When the neurons of a full chain or a large portion of it are observed, the method reveals the associated DS in the intersection matrix. However, when only a few hundreds of neurons (comparable to the number of cells that can be recorded simultaneously *in vivo* with modern electrophysiological techniques) are randomly sampled from the full network, the DSs are less visible and less likely to be continuous, and it is impossible to distinguish them from the background entries. Indeed, this method leaves the judgment whether individual entries are large or small as well as the grouping of proximal events into DSs to visual inspection. As a consequence, the results are prone to subjectiveness and the procedure is not open to automation.

In this chapter we present an improved, automated analysis, named ASSET (Analysis of Sequences of Synchronous EvenTs), which features i) a statistical assessment of membership of individual entries of the intersection matrix to a DS, ii) a rigorous construction of individual DSs and associated SSEs by clustering, and iii) the reconstruction of the neuronal composition and occurrence times of each event in the found SSEs.

The chapter is organized as follows. Building on the notion of the intersection matrix and the idea of filtering, we present in Section 4.2 an extension of the method introduced by Schrader et al. (2008) that determines the membership of matrix entries to a DS by statistical assessment. The assessment is based on two statistical tests for the significance of each individual entry in the matrix and the joint significance of an entry's neighbors, respectively. Entries passing the two tests are grouped together into individual DSs depending on their reciprocal distance by a clustering procedure. Once the DSs are identified, it is readily possible to reconstruct its neuronal composition.

In Section 4.3 we assess the performance and robustness of ASSET on various types of simulated data that replicate typical features of electrophysiological recordings. These include firing rate heterogeneity across neurons, variability over time and different types of correlation structure in the spiking activity. In addition, we demonstrate the ability of the ASSET analysis to find repeated SSE activity in data derived from a neural network model of overlapping SFCs as introduced by Schrader et al. (2008).

We conclude by discussing benefits and limitations of the ASSET method, open problems and future plans for the analysis of electrophysiological data.

4.2 Methods

In this section we provide a formal definition of the intersection matrix as introduced by Schrader et al. (2008), introduce statistical tests to determine whether individual entries of the intersection matrix are part of any DS and propose a clustering technique to group significant entries into the different DSs. After the last step we can determine the occurrence times and neuronal composition of all events of the repeated SSEs associated to the DSs found. We then formally describe the stochastic models employed in Section 4.3 to assess the performance of the method on a variety of scenarios.

4.2.1 The intersection matrix

Our approach builds on the notion of the intersection matrix defined by Schrader et al. (2008). Here we introduce this concept mathematically and provide definitions. Given a set of N parallel spike trains observed in the time interval $[0, T]$, synchronous spike events across neurons are determined by discretizing the time interval into B adjacent time bins b_1, b_2, $\ldots, b_B$ of identical width $\Delta = T/B$ (typically of a few ms), as illustrated in Figure 4.1A. Each set S_i of spikes falling in time bin b_i forms a synchronous event. The value $|S_i \cap S_j|$ represents the number of neurons being active at both time bins b_i and b_j. The *intersection matrix* I is defined by

$$I_{ij} := |S_i \cap S_j| \quad \forall i,\, j = 1, 2, \ldots, B.$$

In this setting I is a symmetric matrix, whose main diagonal contains the population histogram. A synchronous spike event occurring at two bins i and j, e.g. as a result of the repeated activation of the same SFC, results in a larger value for I_{ij} compared to chance level. A repeated SSE composed of l_{SSE} successive synchronous events, each of which occurs at time bins (b_{i_r}, b_{j_r}), $r = 1, 2 \ldots, l_{\mathrm{SSE}}$, determines a sequence of large-valued entries I_{i_r, j_r} in I - which we term a *diagonal structure* (DS) - as sketched in Figure 4.1B.

To account for the variability of firing rates over time when comparing different entries of I, Schrader et al. (2008) normalized each entry of I_{ij} by the number of neurons active in each bin b_i and b_j. After normalization the entries take value 1 for a complete overlap and value 0 for no overlap. To enhance the contrast between high-valued entries belonging to the same diagonal structure and the surrounding entries, the matrix was further filtered by a linear kernel having an orientation parallel to the main diagonal. We here take a different approach as outlined in the following section.

4.2.2 Statistical significance of individual entries in the intersection matrix

To derive a measure of overlap that is independent of firing rates, we first need the probability mass function (pmf) $p_{ij}(\cdot)$ of each individual entry I_{ij} in the intersection matrix, under the null hypothesis H_0 that the spike trains under consideration are mutually independent and Poisson. If the null hypothesis is rejected, i.e. if the observed value ξ taken by I_{ij} is too large to be interpreted as chance, we classify the overlap $S_i \cap S_j$ as a statistically significant repeated synchronous event. In Section 4.3 we show that the statistics of the method are robust to deviations from Poissonianity as well as to the presence of various types of correlations other than repeated SSE activity.

Under the stated hypotheses the distribution of I_{ij} is determined by the firing rate of each neuron k at the time bins b_i and b_j. I_{ij} represents the (stochastic) number of neurons firing simultaneously in both bins and can thus be expressed as

$$I_{ij} = \sum_{k=1}^{N} I_{ij}^{(k)}, \tag{4.1}$$

where $I_{ij}^{(k)}$ is a Bernoulli random variable taking value 1 if neuron k fires in both bins i and j, and 0 otherwise. The probability parameter $p_{ij}^{(k)}$ of $I_{ij}^{(k)}$ is related to the local firing rate $\lambda_i^{(k)}$ of neuron k at bin b_i and $\lambda_j^{(k)}$ at bin b_j by

$$p_{ij}^{(k)} = (1 - e^{-\lambda_i^{(k)}\Delta})(1 - e^{-\lambda_j^{(k)}\Delta})$$

if $i \neq j$, where each factor $1 - e^{-\lambda_i^{(k)}\Delta}$ is the probability for neuron k of having at least one spike at time i, and by

$$p_{ii}^{(k)} = 1 - e^{-\lambda_i^{(k)}\Delta}$$

if $i = j$. The knowledge of the firing rate profiles of all neurons is therefore a prerequisite for the exact computation of the pmf $p_{ij}(\cdot)$ of I_{ij}, $i, j = 1, \ldots, N$.

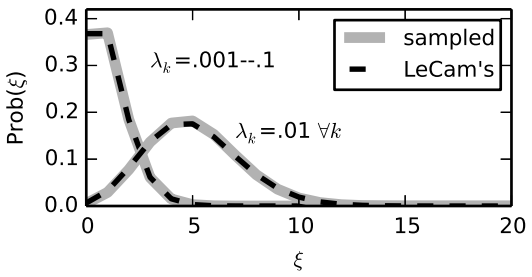

Figure 4.2: **Le Cam's approximation.** The continuous gray curves show two Poisson Bernoulli distributions resulting from the sum of 100 Bernouilli random variables, having rates $\lambda_k = 0.01\ \forall k$ (right curve) and $\lambda_k = 0.001, 0.002, \ldots, 0.1$ (left curve), respectively. The distributions are drawn as histograms from samples of size 10^8 (calculating their exact analytical expression as in Equation (4.2) is computationally prohibitive). In both cases, Le Cam's approximation given by a Poisson distribution with parameter $\lambda = \sum_k \lambda_k$ (dashed lines) yields a good match (integrated absolute error < 0.04).

In light of Equation (4.1), I_{ij} is a Poisson Bernoulli random variable (Le Cam, 1963). Its pmf $p_{ij}(\cdot)$ is analytically given by

$$p_{ij}(\xi) = \sum_{A \in \mathcal{P}_{N;\xi}} \left\{ \prod_{k \in A} p_{ij}^{(k)} \prod_{h \in A^C} (1 - p_{ij}^{(h)}) \right\}, \tag{4.2}$$

where $\mathcal{P}_{N;\xi}$ is the family of all possible subsets of ξ elements that can be extracted from the set $\{1, \ldots, N\}$, A is one such subset, and $A^C = \{1, \ldots, N\} \backslash A$ is the complement of A.

Thus $p_{ij}(\xi)$ is a summation of $\binom{N}{\xi}$ addenda, for a total of 2^N terms needed to compute $p_{ij}(\cdot)$. The computation is feasible for small N, but soon becomes prohibitive as N increases beyond a few dozens, as in the applications to large parallel recordings we are interested in.

By use of Le Cam's theorem (Le Cam, 1963), we approximate $p_{ij}(\cdot)$ by a Poisson density function $p_{(\lambda)}$ with rate parameter $\lambda = \sum_k \lambda_k$

$$p_{ij}(\xi) \simeq p_{(\lambda)}(\xi) = \frac{\lambda^\xi e^{-\lambda}}{\xi!}. \tag{4.3}$$

The approximation error grows quadratically with the $p_{ij}^{(k)}$'s and stays low if the $p_{ij}^{(k)}$'s are sufficiently small, as shown in Figure 4.2.

Given the firing rate profile (i.e. the rate at each time bin) of each neuron, we can thus calculate the pmf $p_{ij}(\cdot)$ and its cumulative distribution function (cdf) $\mathbb{P}_{ij}(\cdot) := \sum_{\xi <\cdot} p_{ij}(\xi)$, either exactly from Equation (4.2) or approximately relying on Le Cam's approximation in Equation (4.3). Transforming each entry I_{ij} by its respective cdf $\mathbb{P}_{ij}(\cdot)$ we can map the observed overlaps I_{ij} to cumulative probabilities

$$P_{ij} := \mathbb{P}_{ij}(I_{ij}) = \sum_{\xi < I_{ij}} \frac{\lambda^\xi e^{-\lambda}}{\xi!}, \tag{4.4}$$

obtaining the *probability matrix* $P := (P_{ij})_{ij}$, as illustrated in Figure 4.3A.

If the null hypothesis holds then $\mathbb{P}_{ij}(\cdot)$ is the true probability distribution of the amount of overlap between bins b_i and b_j, and I_{ij} is a realization from that probability distribution. If so, $P_{ij} = \mathbb{P}_{ij}(I_{ij})$ is uniformly distributed in the unit interval. If P_{ij} is large (close to one) the null hypothesis is rejected in favor of the alternative hypothesis that the observed

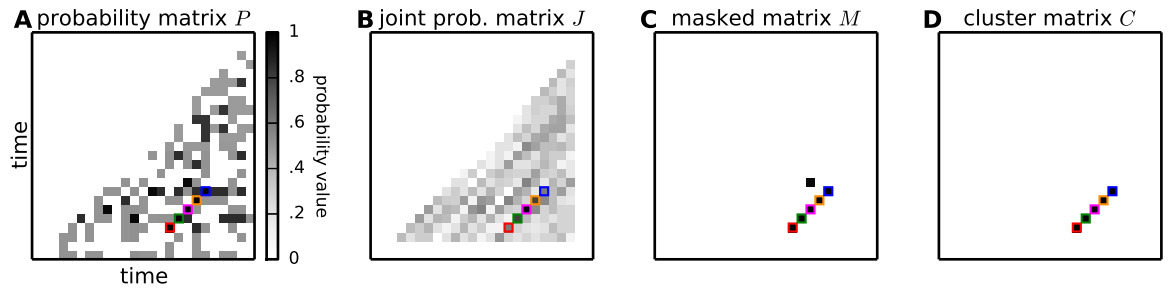

Figure 4.3: **From spike trains to the cluster matrix. (A)** Given parallel spike train data and their intersection matrix I as in Figure 4.1, for each entry I_{ij} its cumulative probability P_{ij} is calculated analytically under the null hypothesis $\mathcal{H}_0$ that the spike trains are independent and marginally Poisson. **(B)** The l largest neighbors of I_{ij} in a rectangular area extending along the 45 degree direction are isolated by means of a kernel and their joint cumulative probability is assigned to the joint probability matrix J at position J_{ij}. **(C)** Chosen a significance threshold α_1 for the probability of individual entries (i.e. for entries P_{ij}) and a significance threshold α_2 for the joint probability of entry's neighbors (i.e. for entries J_{ij}), each entry I_{ij} such that $P_{ij} > \alpha_1$ and $J_{ij} > \alpha_2$ is classified as statistically significant. Significant entries of I are retained in the binary masked matrix M_{ij}, which takes value 1 at positions (i, j) where I is statistically significant and 0 elsewhere. **(D)** 1-valued entries in M falling close-by are clustered together (or discarded as isolated chance events) by means of a DBSCAN algorithm, which thus isolates diagonal structures.

overlap reflects active synchronization between the involved neurons at time bins b_i and b_j. Setting a significance threshold α_1 (e.g. $\alpha_1 = 0.99$) we classify all entries I_{ij} such that $P_{ij} > \alpha_1$ as statistically significant, along with the associated repeated synchronous events.

4.2.3 Joint significance of neighbors of I_{ij}

A DS resulting from a repeated SSE, as sketched in Figure 4.1B, differs from the surrounding entries of the intersection matrix I due to not just one, but a sequence of entries with large values aligned around a diagonal. We devise a statistical test that exploits this feature to detect DSs in the intersection matrix. In the previous section we have already derived the probability matrix P as a transformation of I that normalizes raw intersection values for the local neuronal firing rates. We can now look for DSs in P rather than in I.

Neighbors extracted by rectangular kernel. Schrader et al. (2008) proposed a visual technique to enhance DSs in the intersection matrix by averaging each entry in the matrix with surrounding entries falling in the same diagonal

$$P_{ij} \to F_{ij} = \frac{1}{l} \sum_{h=-l/2}^{l/2} P_{i+h,\, j+h},$$

and thus building a filtered version F of the intersection matrix (in our case, of P). F_{ij} is large if the neighbors $P_{i+h,\, j+h}$ of P_{ij} are jointly large, and small otherwise.

Inspired by this approach we aim at mapping P into a joint probability matrix J whose entries J_{ij} reflect the joint probability of entries surrounding P_{ij}.

We first note that the DS does not have to be composed of adjacent entries in I, and that these entries do not necessarily lie on the same off-diagonal. For instance, if any r-th entry of the DS, $r \in \{1, 2, \ldots, l_{\mathrm{SSE}}\}$, is such that $i_r > i_{r-1} + 1$ and/or $j_r > j_{r-1} + 1$, then

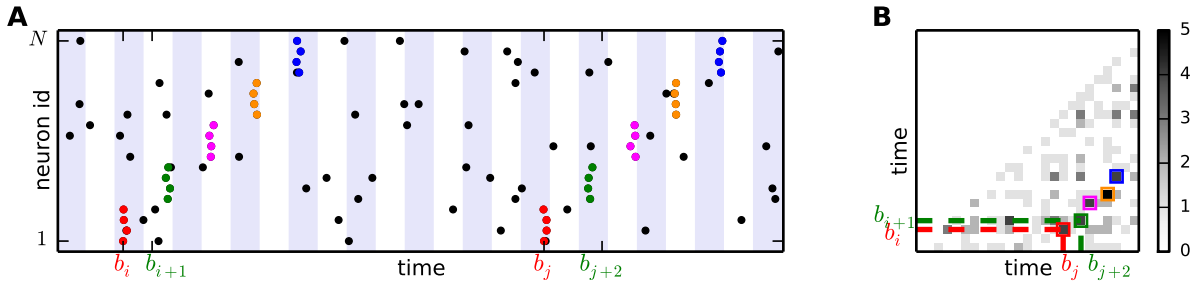

Figure 4.4: **Origin of holes and wiggliness of a DS.** Spiking activity containing a repeated SSE **(A)** and associated intersection matrix **(B)**. The SSE comprises 5 successive events. If two successive events of the SSE do not fall into adjacent bins (e.g. the second and third events in the first SSE repetition) a hole appears in the associated DS. If the distance (in number of bins) of two consecutive events is not identical for the two SSE repetitions, the corresponding entries belong to different off-diagonals of the intersection matrix.

(i_{r-1}, j_{r-1}) and (i_r, j_r) will not be adjacent. If $i_r - i_{r-1} \neq j_r - j_{r-1}$ then (i_{r-1}, j_{r-1}) and (i_r, j_r) will not lie on the same diagonal, as sketched in Figure 4.4. The number of parallel diagonals spanned by the DS, or its "wiggliness", is given by

$$w_{\mathrm{DS}} := \max_r\{\Delta i_r - \Delta j_r\} - \min_r\{\Delta i_r - \Delta j_r\} + 1,$$

where $\Delta i_r := i_r - i_1$ and $\Delta j_r := j_r - j_1$ are the lags (in number of bins) of the r-th synchronous event from the first one, for the first and second occurrence of the SSE, respectively.

To select the entries around P_{ij} for which to calculate the joint probability, we center a rectangular kernel of length l_K and width w_K around P_{ij} , aligned along the diagonal P_{ij} belongs to. If P_{ij} belongs to a DS, the kernel should optimally cover all the entries of the same DS. A rectangular kernel having width $w_K > 1$ allows capturing wiggly DSs, contrarily to a linear kernel. We call the set of matrix entries covered by the kernel the *neighborhood* of P_{ij}. The neighborhood covers $n = l_K \cdot w_K - \left\lfloor \frac{w_K}{2} \right\rfloor \left(\left\lfloor \frac{w_K}{2} \right\rfloor + 1 \right)$ entries in total. Even when P_{ij} belongs to a DS fully covered by the kernel, only some of the n neighbor entries belong to the DS. For this reason we select only the d largest (i.e. most significant) entries in the neighborhood of P_{ij} and evaluate their joint statistical significance. The parameters l_K, w_K and d reflect the putative DS length l_{DS}, DS wiggliness w_{DS} and SSE length l_{SSE}, respectively. Setting ideally $l_K = 2l_{\mathrm{DS}} + 1$ and $w_K = w_{\mathrm{DS}}$ would ensure that the kernel fully covers the DS, also when placed at one end of the DS itself. Setting $d = l_{\mathrm{SSE}}$ would ensure that only the l_{SSE} highest entries (likely to be those forming the DS, if any) are considered, and not all the n entries covered by the kernel. Deviations from these optimal values might lead to consider additional or less entries than those forming the DS. As a rule, we set l_K and w_K odd so that the kernel is symmetric around P_{ij}.

Significance of the d largest kernel entries. We now calculate the joint cumulative distribution of the d out of n largest neighbors of P_{ij} under the null hypothesis H_0 that the spike trains are independent and Poisson. Under H_0 the entries in P are samples from the uniform distribution in the unit interval $F(x) = x$. We additionally consider the d extracted entries as independent samples. This does not hold true if two entries share an index, e.g. (i, j_1) and (i, j_2), but correlations if present are weak and do not affect performance significantly.

For a sample $(X_1, X_2, \ldots, X_n)$ of n independent realizations extracted from a probability distribution F defined in the interval $[0,1]$, the joint survival function $\overline{F}$ of the d largest order statistics $X^{(n-d+1)} < X^{(n-d+2)} < \ldots < X^{(n)}$ extracted from the sample is defined by

$$\overline{F}_{n-d+1,n-d+2,\ldots,n}(x_1, x_2, \ldots, x_d) := Pr(X^{(n-d+1)} \geq x_1,\ X^{(n-d+2)} \geq x_2, \ldots,\ X^{(n)} \geq x_d)$$

and represents the probability that at least d samples are larger than or equal to x_1, at least $d-1$ samples are larger than or equal to x_2, ..., and at least 1 sample is larger than or equal to x_d.

Since the $x_1 \leq \ldots \leq x_d$ are ordered, the distribution vanishes if at least d samples are not $\geq x_1$. Hence the number i_1 of samples $\geq x_1$ can take values $i_1 = d, d+1, \ldots, n$ and the remaining $n - i_1$ samples are below x_1. Of the i_1 samples $\geq x_1$, there must be at least $d-1$ samples $\geq x_2$. Of the i_2 samples $\geq x_2$, $i_2 = d-1, d, \ldots, i_1$, there must be at least $d-2$ samples $\geq x_3$ and so on. Continuing this reasoning we arrive at the expression

$$\overline{F}_{n-d+1,n-d+2,\ldots,n}(x_1, x_2, \ldots, x_d) = \sum_{i_1=d}^{n} \sum_{i_2=d-1}^{i_1} \ldots \sum_{i_{d-1}=2}^{i_{d-2}} \sum_{i_d=1}^{i_{d-1}} n! \prod_{k=0}^{d} \frac{[F(x_{k+1}) - F(x_k)]^{i_k - i_{k+1}}}{(i_k - i_{k+1})!},$$

for any $x_1 \leq x_2 \leq \ldots \leq x_d$, and 0 otherwise, where $x_0 = 0$, $x_{d+1} = 1$, $i_0 = n$ and $i_{d+1} = 0$.

Substituting to F the identity function $F(\cdot):\ F(x) = x$, to x_{d-k+1} the largest values $P^{(d-k+1)}$ in the neighborhood of P_{ij} and denoting $\overline{F}_{n-d+1,\ldots,n}$ by $\overline{F}$ for simplicity, yields

$$\overline{F}(P^{(1)}, P^{(2)}, \ldots, P^{(d)}) = n! \sum_{i_1=d}^{n} \ldots \sum_{i_{d-1}=2}^{i_{d-2}} \sum_{i_d=1}^{i_{d-1}} \prod_{k=0}^{d} \frac{(P^{(k+1)} - P^{(k)})^{i_k - i_{k+1}}}{(i_k - i_{k+1})!},$$

which represents the tail probability (test p-value) of the joint event

$$X^{(n-d+1)} \geq P^{(1)} \wedge\ X^{(n-d+2)} \geq P^{(2)} \wedge\ \ldots \wedge\ X^{(n)} \geq P^{(d)}$$

under H_0. The p-value is small (that is, reflects high significance) either if $P^{(1)}, P^{(2)}, \ldots, P^{(d)}$ are jointly large - i.e. if the corresponding intersection values I_{ij} are statistically large - or if just one of them (e.g. the largest, $P^{(d)}$) is very large. To avoid the second scenario we impose an upper bound on the values of P by replacing each $P^{(r)}$, $r = 1, \ldots, d$, with $P_*^{(r)} = \min(P^{(r)}, p_{\max})$ (e.g. $p_{\max} = 0.999$) to obtain

$$\overline{F}(P^{(1)}, P^{(2)}, \ldots, P^{(d)}) = n! \sum_{i_1=d}^{n} \ldots \sum_{i_{d-1}=2}^{i_{d-2}} \sum_{i_{d-1}=1}^{i_{d-1}} \prod_{k=0}^{d} \frac{(P_*^{(k+1)} - P_*^{(k)})^{i_k - i_{k+1}}}{(i_k - i_{k+1})!}. \qquad (4.5)$$

By doing so, statistically highly significant values of $\overline{F}$ are possible only in the presence of jointly large neighbors of P_{ij}.

The complementary function $1 - \overline{F}(\cdot)$ returns the probability of *not* having the joint event, and can be used to map the probability matrix P into a *joint probability matrix* J

$$P_{ij} \rightarrow J_{ij} = 1 - \overline{F}(P^{(1)}, P^{(2)}, \ldots, P^{(d)}). \qquad (4.6)$$

Figure 4.3B illustrates the joint probability matrix J derived from the probability matrix P in panel A.

Setting a significance threshold α_2 (e.g. $\alpha_2 = 0.99999$) we classify entries I_{ij} in the raw intersection matrix as having significantly jointly large neighbors if $J_{ij} > \alpha_2$.

4.2.4 Clustering entries of I into DSs

Each entry I_{ij} in the raw intersection matrix is tested for individual significance and joint significance of its neighbors. If both tests pass, i.e. if $P_{ij} > \alpha_1$ and $J_{ij} > \alpha_2$, the entry is classified as belonging to one DS. Such entries are collected in a binary *masked matrix* M that takes value 1 if both tests pass, and 0 otherwise

$$M_{ij} := 1_{\{I_{ij}>\alpha 1\}} \cdot 1_{\{J_{ij}>\alpha_2\}}, \tag{4.7}$$

as illustrated in Figure 4.3C. It remains to be established which entries belong together to the same DS. Intuitively, 1-valued entries in the masked matrix belong to the same DS if close-by, and to different DSs if far apart. The masked matrix in Figure 4.3C shows for instance two clearly separated DSs.

A suitable notion of "distance" for matrix entries should make entries falling in the same diagonal, i.e. aligned along the natural direction of a DS, closer together than entries aligned along the anti-diagonal. We introduce the following elliptic distance between any two matrix entries (i_1, j_1) and (i_2, j_2)

$$d_\rho((i_1, j_1), (i_2, j_2)) := \left[1 + (\rho - 1) \cdot \sin\left(\left|\theta - \frac{\pi}{4}\right|\right)\right] \cdot \sqrt{(i_2 - i_1)^2 + (j_2 - j_1)^2}/\sqrt{2},$$

where $\theta = \arctan\left(\frac{j_2-j_1}{i_2-i_1}\right)$ is the angular coefficient of the line intersecting (i_1, j_1) and (i_2, j_2), the first square root factor is the Euclidean distance between the two points and $\rho \geq 1$ is a stretching factor for angular coefficients deviating from $\pi/4$. $d(\cdot)$ grows as θ approaches $3\pi/4$ or $-\pi/4$, i.e. the anti-diagonal orientation. For instance, $d_\rho((i, j), (i+k, j+k)) = k$ for any ρ, $d_\rho((i, j), (i+k, j-k)) = \rho k$ and $d_\rho((i, j), (i, j+k)) = \left[1 + \frac{\sqrt{2}}{2}(\rho - 1)\right] k$.

We set $\rho = 10$ and, based on the distance $d_{10}(\cdot, \cdot)$ of their positions in the matrix, group all entries I_{ij} with $M_{ij} = 1$ (see Equation (4.7)) into clusters via a density-based scanning (DBSCAN) algorithm (Ester et al., 1996). The algorithm considers two entries as part of the same neighborhood if their distance is not larger than a maximum value ε. Neighborhoods sharing an entry are joint together and eventually classified as a cluster if they contain a minimum number l_0 of entries. We set $\varepsilon = 5$, thus allowing for a maximum of $5 - 1 = 4$ holes between two consecutive entries of a DS along the main diagonal, and $l_0 = 3$, thus requiring a DS to reflect at least 3 repeated synchronous events. Entries in the matrix M not belonging to any cluster are discarded as events that do not reflect repeated SSE activity.

Figure 4.3D shows the *cluster matrix* C assigning value 1 (colored in black) to entries belonging to a cluster, and 0 (white) to the others. The matrix contains two clusters composed of 5 entries each, corresponding to the 5 synchronous events highlighted in Figure 4.1A.

4.2.5 Rate estimation methods

Calculating the entries of the probability matrix as given by Equation (4.2) and Equation (4.3) requires the knowledge of the firing rate of each neuron over time. Firing rate estimation in a problem that has been targeted by a number of studies.

The peri-stimulus time histogram (PSTH, Perkel et al., 1967a) is an estimate of the firing rate performed by discretizing time into adjacent bins and by computing the number of spikes falling into each bin. The bin width for the PSTH is typically larger (tens to hundreds of ms) than the one used to define synchrony, as firing rates change on a slower time scale. The larger the bin width, the coarser but less biased the estimate. Normalizing

the spike count in each bin by the bin width yields the rate of the process, i.e. the number of spikes per time unit.

Kernel convolution (Nawrot et al., 1999) replaces each spike with a kernel (a probability density function) centered around the spike time and estimates the firing rate by the sum of these distributions. Formally, this is done by representing the spike time by a Dirac delta function centered around the spike time t^* and by convolving it with the kernel. Following a standard choice, we specify the kernel as a normal distribution with assigned standard deviation σ, truncated at $\pm 2.7\sigma$ to yield a finite support. When setting the kernel width $w^* = 5.4\sigma$ to a fixed value, we employ $w^* = 200\,\mathrm{ms}$.

Both PSTH and kernel convolution can be applied to the case when multiple independent, identically distributed trials of the activity of a neuron are available, by averaging the estimates obtained for each trial. For identical bin and kernel widths, the PSTH typically better represents sharp changes in the firing rate from one bin to the next, while kernel convolution yields smoother curves.

Both estimates are parametric and require the choice of a bin- (kernel-) width. Methods have been recently proposed to determine the optimal bin- or kernel-width by minimizing the error between the true (unknown) rate and its estimate in some statistical sense (see Shimazaki and Shinomoto, 2007, 2010). These methods have been shown to outperform their fixed-width variants and are particularly helpful when analyzing parallel spike train data from different neurons with different rates, where the optimal bin width varies across neurons.

In Section 4.3 we compare the performance of ASSET employing PSTH, kernel convolution, and optimized-width kernel convolution (Shimazaki and Shinomoto, 2010) estimates of the firing rate profiles over an increasing number of trials.

4.2.6 Stochastic models of background activity

We devise 10 different modes of background spiking activity embedding and combining different features typical of experimental data. We provide here the definition of each model. Examples of the spiking activities of each model can be seen in Figure 4.6.

Model 0 - independent Poisson. $N = 100$ independent Poisson spike trains which have identical, stationary firing rate $\lambda = 15\,\mathrm{Hz}$ each.

Model 1 - Poisson with simultaneous rate jump. $N = 100$ Poisson spike trains having an identical, time-varying firing rate profile $\lambda(t) = 10\,\mathrm{Hz} + 50\,\mathrm{Hz} \cdot 1_{\{600\,\mathrm{ms}<t<700\,\mathrm{ms}\}}$. Thus, all spike trains undergo a simultaneous sudden rate excursion of $+50\,\mathrm{Hz}$ between times $t_1 = 600\,\mathrm{ms}$ and $t_2 = 700\,\mathrm{ms}$. The average rate over the 1 s simulation period is $15\,\mathrm{Hz}$, as in model 0.

Model 2 - heterogeneous, time-stationary Poisson. $N = 100$ independent Poisson spike trains with heterogeneous firing rates across neurons ranging from $5\,\mathrm{Hz}$ to $25\,\mathrm{Hz}$. Neuron k has stationary firing rate $\lambda_k = 5 + \frac{20}{99}k\,\mathrm{Hz}$, $k = 0, 1, \ldots, 99$. The population-averaged firing rate is $15\,\mathrm{Hz}$, identical to models 0 and 1.

Model 3 / 4 / 5 - high inter-spike interval regularity. Same as models 0 / 1 / 2 respectively, but the spike trains have marginally Gamma distributed inter-spike intervals (ISIs) with shape factor $\alpha = 5$ ($\gamma = 1$ for the Poisson case) and mean parameter μ

$$\mathbb{P}(\text{ISI} = \tau) = \left(\frac{\alpha}{\mu}\right)^{\alpha} \frac{\tau^{\alpha-1} e^{-\alpha\tau/\mu}}{\Gamma(\alpha)}.$$

The Gamma distribution determines ISIs having mean length μ (firing rate: $1/\mu$), standard deviation $\sigma = \mu/\sqrt{\alpha}$ and thus a coefficient of variation $\sigma/\mu = 1/\sqrt{\alpha}$. The larger α, the more regular are the ISIs compared to their mean value.

Model 6 - fast rate-jump propagation. $N = 100$ independent Poisson spike trains, organized in 20 groups of 5 spike trains each. The first group experiences a sudden, simultaneous rate jump from $\lambda_1 = 14\,\text{Hz}$ to $\lambda_2 = 100\,\text{Hz}$ at two times $t_1 = 50\,\text{ms}$ and $t_2 = 500\,\text{ms}$, lasting 5 ms (one time bin) each time. Then the second group undergoes the same rate jump 5 ms (one bin) later, and so on. The l-th group witnesses the rate jump at time $t_1 + 5(l-1)\,\text{ms}$ and $t_2 + 5(l-1)\,\text{ms}$, $l = 1, \ldots, 20$.

Model 7 - population synchronization. $N = 100$ spike trains having all stationary firing rates $\lambda = 15\,\text{Hz}$. From time to time, 5 randomly selected neurons fire synchronous spikes, with a frequency that yields a mean pairwise correlation coefficient among any two spike trains of $\rho = 0.01$. The rest of the spiking activity is mutually independent.

Formally, the model is described by a compound Poisson process (CPP; Staude et al., 2010b,a) with a two-peak amplitude distribution $A(\cdot)$ such that $A(1) = 0.938$, $A(5) = 0.062$ and $A(\xi) = 0\ \forall \xi \neq 1, 5$. Each neuron has an average participation rate in synchronous events of 4.63 Hz.

Model 8 - intra-groups synchronization. Multiple Single-Interaction Process (mSIP, see Grün et al., 2008; Gerstein et al., 2012; cf. Kuhn et al., 2003b for details on the basic SIP model) comprising 65 independent neurons plus 7 groups (SIPs) of 5 neurons each, for a total of $N = 100$ neurons. Each SIP exhibits synchronous activity in two randomly selected time bins, independent for each group. Thus, the resulting data contains 7 disjoint repeated synchronous events of size 5. This is the same as for the case when we inject a repeated SSE into the data, with the difference that the synchronous events do not form a temporal sequence here. Each neuron has a total firing rate $\lambda = 15\,\text{Hz}$, stationary over time.

Model 9 - time-varying heterogeneous Poisson. $N = 100$ independent, marginally Poisson spike trains whose firing rates are heterogeneous across neurons and non-stationary over time. Spike train k has rate profile

$$\lambda_k(t) = \left(5 + \frac{10 \cdot k}{99}\right)\text{Hz} + 50\,\text{Hz} \cdot 1_{\{600\,\text{ms} < t < 700\,\text{ms}\}},$$

$k = 0, 1, \ldots, 99$. Thus, the spike trains have baseline firing rates ranging from 5 Hz to 15 Hz and witness a coherent rate jump of +50 Hz between 600 ms and 700 ms. The time- and population-averaged firing rate is 15 Hz. This model mixes non-stationarity and heterogeneity of firing rates individually characterizing models 1 and 2, respectively.

4.3 Results

We propose here an extension of the method presented by Schrader et al. (2008) that visualizes repeated temporal sequences of synchronous events (SSE) as diagonal structures (DS, a sequence of large values along a diagonal) within an intersection matrix I. I_{ij}

represents the number of neurons which have spikes both in time bin i and time bin j (see Figure 4.1 and Equation (4.1) for the formal definition). We map the intersection matrix I into a probability matrix P which contains at each position P_{ij} the probability of an overlap lower than I_{ij} between the spike trains at the two time bins (Figure 4.3A). The calculation is derived analytically under the null hypothesis H_0 that the spike trains are Poisson and independent. We then further compute the joint probability matrix J whose entries J_{ij} represent the cumulative joint probability of neighbors of I_{ij}, again under H_0 (Figure 4.3B). Set two statistical thresholds $\alpha_1 < \alpha_2$, each entry such that both $P_{ij} > \alpha_1$ and $J_{ij} > \alpha_2$ hold is classified as a potential member of a DS. Technically, this is done by building a masked matrix M such that $M_{ij} = 1$ if the two conditions hold, and $M_{ij} = 0$ otherwise (Figure 4.3C). 1-valued entries in M are finally clustered into individual DSs or discarded as isolated entries based on their reciprocal distance, thus yielding a cluster matrix C (Figure 4.3D).

The positions (i, j) of entries composing each DS, which can be obtained from matrix C, allow us to reconstruct the synchronous events forming the associated repeated SSE as the intersection of the neurons active at bins i and j.

We investigate the performance of the proposed method on simulated test data. The data are generated by stochastic simulations of N parallel spike trains (here we chose $N = 100$) over a time period of 1 s. We define 10 types of background spiking activity, differing by the marginal properties of the spike trains (firing rate profile and ISI distribution) and the correlations among the spike trains. We use these data to determine the false positive (FP) rate of the method in cases where SSEs are not included in the data. Then we enrich the data with two occurrences of an SSE composed of $l_{\mathrm{SSE}} = 7$ synchronous events, each one involving $\xi_{\mathrm{SSE}} = 5$ neurons with IDs $1-5$, $6-10$, ..., $31-35$, and investigate therein the true positive (TP) and FP rates. In data containing SSE activity a found SSE is considered as an FP if it is not composed of (or it is only partly composed of) the events forming the embedded SSE. Details are given below. For the diversity of data types we test the power of the method , and further identify critical cases and suggest solutions.

4.3.1 Features and models of stochastic test data

To assess the quality of the method, we measure its performance in the case when all assumptions entering the derivation of Equation (4.2) (alternatively Equation (4.3)) and Equation (4.5) are met. In addition, we test the robustness of the statistics of the method with respect to deviations from these assumptions which are typically found in experimental data. In particular, we investigate how the following features of the data affect the performance of the method and test if these lead to false positive (FP) outcomes. The first four features relate to aspects of firing rates, such as various types of non-stationarities and rate correlations, i.e. correlations between the spike trains on a slower time scale than SSEs. The last two features relate to spike synchrony, however with a different organization that in SSEs.

A. Variability of firing rates over time by means of a sudden rate jump that is coherent across all neurons (Figure 4.5A).
   Firing rate changes over time are the basic observation of experimental data, in particular in response to an external stimulus or in relation to behavior. Neurons reacting to a common stimulus often exhibit a co-modulation of their rates (*noise correlation*, [REFS]). Ignoring or mis-estimating such rate changes is a typical generator of FPs in synchrony analyses (Grün et al., 2003; Grün, 2009; Grün and Rotter, 2010). We thus investigate a worst case scenario of coherent and instantaneous rate changes.

B. Heterogeneity of the firing rates across neurons (Figure 4.5B).
The firing rates of simultaneously observed neurons typically differ. Although we estimate the firing rates on a neuron by neuron basis, we still want to ensure that the method can cope with this rate variability. Such variability combined with cross-trial variability was shown also as a strong generator of FPs (Grün, 2009; Grün and Rotter, 2010, see e.g.).

C. Increased regularity of the spiking activity compared to the Poisson assumption, by means of Gamma-distributed ISIs (Figure 4.5C).
Deviations from Poisson statistics in terms of ISI regularity, which is observed in experimental data (Shinomoto et al., 1999; Farkhooi et al., 2009, 2011), was shown to bias synchrony analysis methods (Grün, 2009; Pipa et al., 2013a) that, like ASSET, assume Poisson spike trains in the null hypothesis.

D. Short lasting, simultaneous rate jumps in a group of neurons, propagating to other groups in a sequence; this "chain" of rate jumps re-occurs at a later time (Figure 4.5D).
This model mimics a rate propagation model, instead of spike synchrony propagation as represented by SSEs. In the ongoing debate on rate coding versus temporal coding (Shadlen and Newsome, 1994; Softky, 1995; Shadlen and Newsome, 1995) it was proposed that coherent short lasting firing rate changes at the input of neurons would be as efficient in bringing neurons to emit a spike as synchronous input. However, rate correlation and spike synchrony are mathematically different constructs and can be differentiated by correlation analysis (Staude et al., 2008). We aim to ensure that ASSET does not misinterpret sequences of coherent rate changes as SSEs, which it is designed to detect.

E. Population synchronization, represented by the occurrence of synchronous spike events at random time points, each involving a random selection of neurons (Figure 4.5E).
ASSET operates under the null hypothesis of spike train independence given the estimated firing rate of each neuron. Thus, fine temporal correlations are not incorporated in the null hypothesis. However, studies of recurrent neuronal networks show the presence of weak temporal correlations (Helias et al., 2013), caused e.g. by the recurrent connectivity in the network. To study their effect on the performance of our method, we generate test data that contain population correlations with unspecific neuronal compositions. As a correlation model we choose the compound Poisson process, which inserts synchronous spikes at a predefined occurrence rate in randomly selected sets of neurons, rather than in specific groups of neurons as for the SSEs.

F. Intra-groups synchronization, where multiple disjoint groups of neurons exhibit synchronous spiking within the group, at independent points in time for each group (Figure 4.5F).
This model represents a second type of correlation structure at fine temporal scale differing from SSE activity. Here, synchronous spiking activities of specific groups of neurons exist in the same number (7) and size (5) as for the embedded SSE, but in contrast to the SSEs are not activated in a temporal sequence. This type of correlation was explored already by Gerstein et al. (2012), were it was shows to produce in the intersection matrix isolated high-valued entries , but no DSs. .We test here whether this holds for our more advance method.

For the concrete test cases we formulate 10 different stochastic models of spiking activity (see Figure 4.6 as examples of the realizations), each including only one or a combination

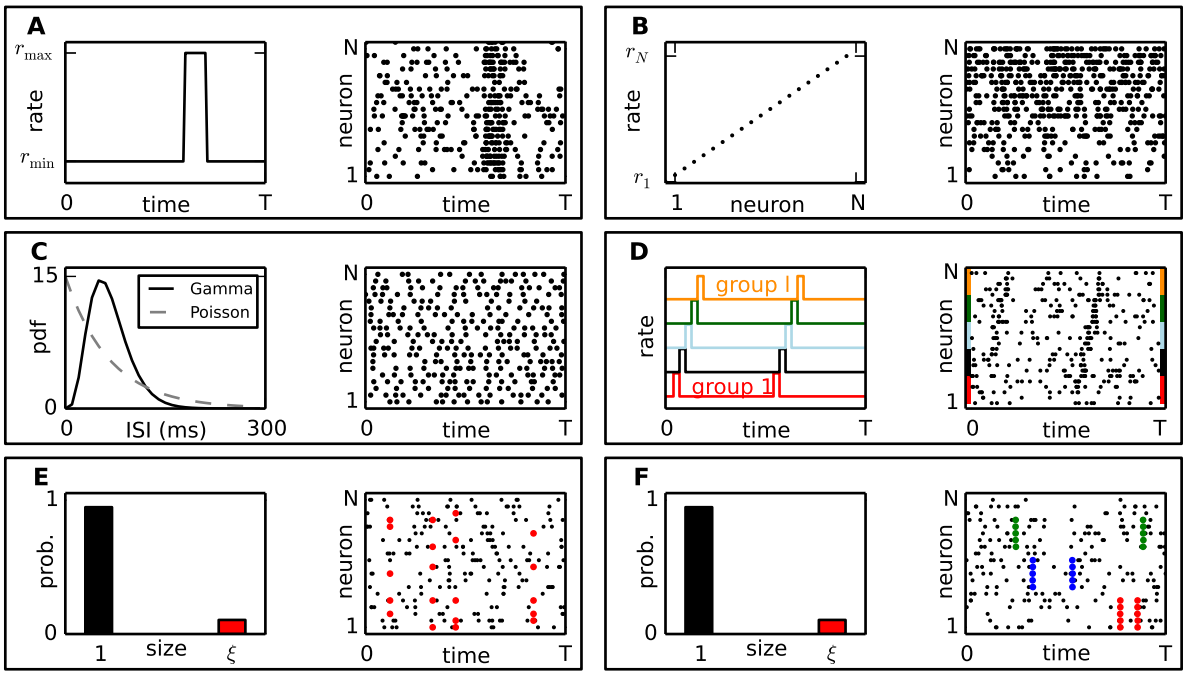

Figure 4.5: **Different features of the data.** Each box illustrates one feature (from A to F) of the test data (left column) and an associated exemplifying raster plot (right column).The dots in the raster plots mark the occurrence time of each spike. **(A)** Variability of firing rates over time. The underlying rate profile is shown on the left. **(B)** Heterogeneity of the firing rates across neurons: the different stationary firing rate levels are marked on the left. **(C)** Gamma-distributed ISIs: the underlying Gamma distribution (solid line) is contrasted by to a Poisson distribution (dashed). **(D)** Short lasting, sequential rate jumps (different colors) coherent within groups of neurons, repeat identically at two time points. **(E)** Population synchronization: the occurrence probability of the size of the synchronous, but randomly selected spike events, is shown in red, the occurrence probability of independent background spikes (size 1) in black. **(F)** Intra-groups synchronization: the complexity distribution is identical to the case of population synchronization in panel E, but the events involve specific groups of neurons (each group is marked in the dot display by one color).

	model id									
	0	1	2	3	4	5	6	7	8	9
A: firing rates varying over time		x			x					x
B: heterogeneous rates across neurons			x			x				x
C: high regularity of ISIs				x	x	x				
D: rate propagation							x			
E: population synchronization								x		
F: intra-groups synchronization									x	

Table 4.1: **Features embedded in the stochastic models.** The crosses mark the features (A to F) included and combined in each model (0 to 9) to simulate the background activity of the test data.

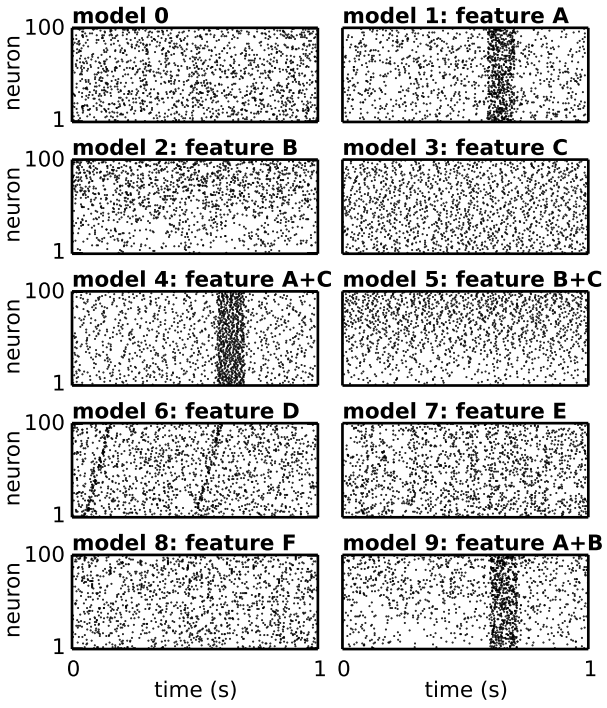

Figure 4.6: **Spiking activity of the stochastic models.** Population raster plots of the test data generated by each model $0-9$. Each panel shows with black dots the spike times of 100 neurons (vertical axis) over time (horizontal axis). Model 0 consists of independent Poisson spike trains with firing rates that are stationary over time and identical across neurons. The other models destroy one or more of these aspects by including some of the features A to F, as listed in Table 4.1. A formal definition of each model is provided in Section 4.2.6.

of the above-mentioned features, as summarized in Table 4.1. We use these models to generate background activity into which we subsequently embed the spiking activities of an repeated SSE. The formal definition and the parameters of each model can be found in Section 4.2.

Figure 4.6 shows example raster plots of the activity associated to each of these stochastic models, without additional injection of SSE activity.

4.3.2 Calibration on test data

p-values of the two statistical tests. Equation (4.3) provides an analytical expression of the cumulative probability P_{ij} of the intersection value I_{ij} under the null hypothesis H_0 of independent, marginally Poisson spike trains. The complementary test p-value $1-P_{ij}$ represents the probability that the intersection between bins b_i and b_j is larger than or equal to I_{ij}. Analogously, Equation (4.5) provides the joint p-value $1-J_{ij}$ of neighbors of P_{ij}. For positions (i,j) of the intersection matrix corresponding to a real DS, J_{ij}

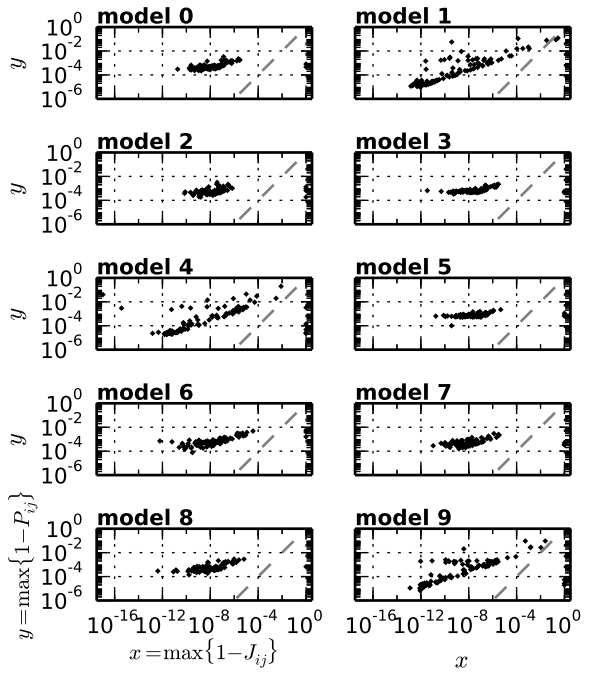

Figure 4.7: **(Scatter of the p-values of the two tests in the presence of a repeated SSE.** The p-value of the first test (y-axis) in relation to the p-value of the second test (x-axis) for all models $0, \ldots, 9$. Each model is repeated 100 times and in each repetition an SSE composed of 7 links and 5 neurons/link is injected at two random points in time. The represented p-values $y = \max_k(1 - P_{i_k,j_k})$ and $x = \max_k(1 - J_{i_k,j_k})$ are taken as the maximum (i.e. least significant) p-value of all entries $\mathcal{S}_{DS} := \{(i_k, j_k),\, k = 1, 2, \ldots, 7\}$ of the intersection matrix composing the embedded DS. The figure shows for each model the cloud of the 100 points $(x,\, y)$, one per simulation.

should take values which are orders of magnitude lower than P_{ij}, as it represents the joint significance of multiple rather than individual repeated synchronous events. To confirm this, we simulate data from each background activity model 0 to 9 (see Figure 4.6) for 100 times. For each simulation, we additionally inject at two points in time, chosen at random, an SSE composed of 7 links and 5 neurons/link, resulting in a real DS composed of a set of 7 entries $\mathcal{S}_{DS} := \{(i_k, j_k),\, k = 1, 2, \ldots, 7\}$ in the intersection matrix. We estimate the firing rate profile of each neuron from single spike trains by kernel convolution (boxcar kernel, kernel width: 200 ms), segment the data in 5 ms bins and calculate the matrices P and J. Given the large number of neurons involved, we need to rely on Le-Cam's approximation of the pmf of intersection values (see Equation (4.3)) to compute P. We then extract from each matrix the p-values $1-P_{i_k,j_k}$ and $1-J_{i_k,j_k}$ of the entries forming the embedded DS, and consider the largest (i.e. least significant) ones, $y = \max_k(1 - P_{i_k,j_k})$ and $x = \max_k(1 - J_{i_k,j_k})$. Figure 4.7 shows for each model a cloud of the 100 points (x, y), one per simulation. As expected, for entries belonging to the true DS the joint significance J_{ij} is orders of magnitude higher than the individual significance P_{ij}. For this reason it makes sense to set the respective statistical thresholds α_2 and α_1 such that $\alpha_2 > \alpha_1$. In particular, on the basis of the scatter plots in Figure 4.7, which show that most p-values P_{i_k,j_k} and J_{i_k,j_k} are lower than 10^{-2} and 10^{-5} respectively, we set $\alpha_1 = 10^{-2}$ and $\alpha_2 = 10^{-5}$.

True positive and false positive DSs. For each simulation of models 0 to 9 we identify statistically significant entries (i, j) in the intersection matrix. We then cluster these entries into diagonal structures using the DBSCAN algorithm, as explained in Section 4.2. Table 4.2 summarizes the parameter values used in the analysis.

We define a DS found in stochastic simulations containing SSE activity as a TP if it contains at least 50 of the entries of the true DS, and if at least 50 of its entries belong to the true DS. If either requirement is not met, or if the data does not contain embedded

intersection matrix		statistical assessments		DBSCAN clustering	
parameter	value	parameter	value	paramete	value
bin size Δ	5 ms	1^{st} thresh. α_1	10^{-2}	stretch coeff. ρ	5
kernel length l	5	2^{nd} thresh. α_2	10^{-5}	max distance ε	3.5
kernel width w	5			min cluster size s	3
# largest neighbors d	5				

Table 4.2: **Analysis parameters.** Parameters of the ASSET method employed for the analysis of stochastic and network data. Each column shows the parameters employed for the corresponding step of the method.

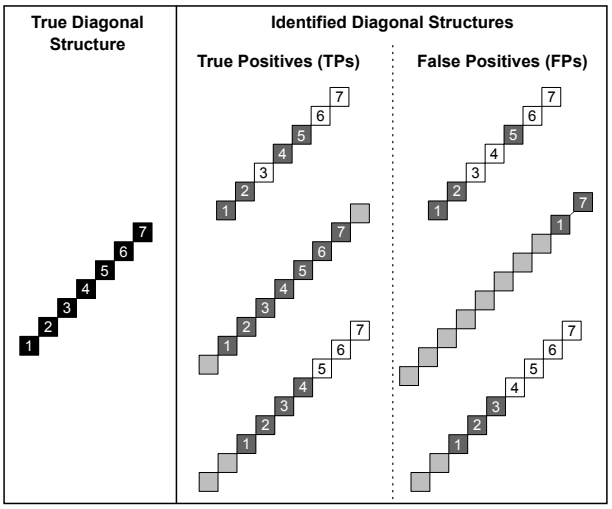

Figure 4.8: **Illustration of true positive and false positive DSs. Left:** DS composed of 7 entries, resulting from the two occurrences of the SSE embedded in the test data. **Middle, right:** A found DS in the cluster matrix can contain entries that belong to the true DS (dark gray) or that do not (light gray), and may miss entries of the true DS (white). **Middle column:** Different types of TP DSs found in data: contained in (*top*), containing (*middle*) and partially overlapping with (*bottom*) the true DS. **Right column:** Different types of FP DSs found in data: contained in (*top*), containing (*middle*) and partially overlapping with (*bottom*) the true DS.

SSE activity in the first place, the found DS is classified as a FP. Figure 4.8, left, illustrates the entries composing a true DS in data with embedded SSE activity. Figure 4.8, right, illustrates FPs arising from the fact that the found DS contains less than 50% of the entries composing the true DS (top, first requirement not met), or contains more than 50% of entries not belonging to the real DS (middle, second requirement not met). The central column in the figure similarly illustrates different types of TPs, when both criteria are satisfied. Note that both TP and FP DSs can be contained in the real DS (top), contain it (middle) or partially overlap with it (bottom).

Figure 4.9, *top*, shows the FP rate (i.e. average number of FPs over the 100 simulations) in each test data model introduced above, for the case where no SSE is embedded in the data (therefore, each found DS is a FP). The middle and bottom panels show the TP and FP rate for the corresponding cases where two occurrences of an SSE are embedded in the data. Optimal performance is achieved in SSE-free data when the FP rate, which can take values between 0 and $+\infty$, is 0. In data with embedded SSEs the TP rate, ranging between 0 and 1, is optimally 1.

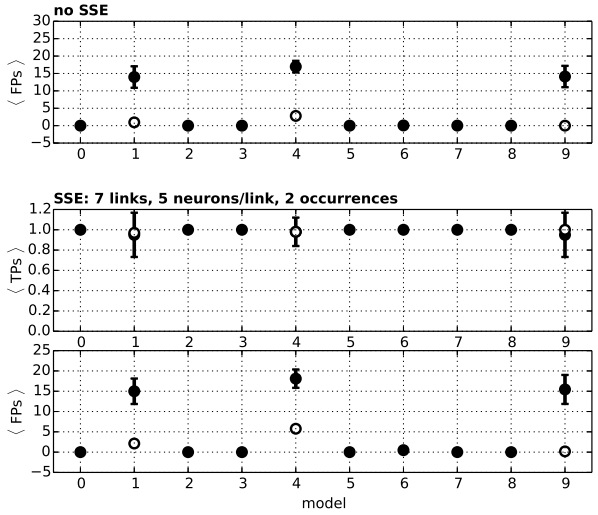

Figure 4.9: **Performance of the method for different models of background activity.** The 10 different models of background (SSE-free) activity, numbered from 0 to 9, are simulated 100 times each. **Top panel:** Average number of FPs found in the data when estimating the firing rates of individual spike trains by kernel convolution (filled circles) and when using the theoretical rate profiles that underlie the generation of the test data (empty circles). Bars indicate the standard deviation. **Center and bottom panel:** Average number of TPs and FPs found in the data containing 2 repetitions of an SSE of $l_{\mathrm{SSE}} = 7$ synchronous events involving $\xi_{\mathrm{SSE}} = 5$ neurons each.

Assumptions of null hypothesis met. When the null hypothesis of Poisson and independent spike trains holds, and the firing rates are stationary over time (model 0, identical rates across neurons, and model 2, different rates) the performance is high both in terms of TP rate (= 1.00) and FP rate (= 0.00, both in test data with and without embedded SSEs). Cross-neuron heterogeneity of firing rates does not affect the performance. These results demonstrate the adequacy of kernel convolution, used to estimate the rate profiles on a single spike train basis, when the rates are stationary. The high performance also indicates the efficacy of Le Cam's approximation (Equation (4.3)) of the true but intractable probability distribution of intersection values (Equation (4.2)).

ISI regularity. The optimal TP and FP levels observed in model 0 and 2 are maintained when the spike trains exhibit higher ISI regularity than Poisson (model 3 and 5, respectively). This shows that the method is robust to deviations from the Poisson assumption employed to derive the null distribution.

Rate propagation. We further evaluate whether high and short-lasting rate increases affecting a group of neurons simultaneously and propagating from one group to the next in a sequence would be interpreted by the method as the occurrence of an SSE. The main difference between rate propagation and SSE activity is that the first causes neurons to fire stochastically (probability < 1) rather than reliably (probability $= 1$). The first model converges to the second as the rate increase grows to infinity, so that ASSET should find SSE activity in this case. It is otherwise understood as a different model of information coding, which should not yield SSEs.

Using model 6, we simulate a border-line scenario in which each of 20 groups of neurons successively increases its firing rate from a baseline of 14 Hz to 100 Hz for a period of 5 ms, with an inter-group delay of 5 ms matching the analysis bin width. This fast rate change propagation occurs twice, thus resembling a repeated SSE (see Figure 4.5D and Figure 4.6, model 6), but with the firing probability of each neuron of 0.39 during the high-rate regime (instead of 1 for SSE activity and 0.067 during the baseline regime). With these parameters the method yields a FP rate of 0.48 (equivalent to less than 1 FP every 2 model iterations),

despite the unrealistically high and sudden rate jump, its coherence across neurons and the fact that the high-rate states perfectly falls into adjacent analysis bins. Deviations from these three features would decrease the discovery rate substantially. On the contrary, even higher rate jumps would increase it. This indicates that the method distinguishes rate propagationfrom SSE activity as long as the two models are substantially different, but would identify the two when the two models converge.

Correlated spike trains. The presence of synchrony involving different groups of neurons each time (model 7) does not harm the performance of the method: the TP and FP rates stay at levels 1.00 and 0.00, respectively. The same results are obtained in data characterized by specific groups of neurons that synchronize their activity within the group, however independently among groups (model 8). These results taken together indicate that, for the range of parameters employed here, the presence of synchronous activity, at the population level or even involving specific groups of neurons, does not affect the performance of the method.

In general, however, one may explicitly include such correlations in the null-hypothesis. To this end we suggest here a Monte-Carlo approach to estimate the probability matrix J while taking the observed repeated synchronous events into account. The basic idea is to estimate J from surrogates obtained by manipulations of the intersection matrix I by destroying the arrangement of rows and columns of I. Using this approach, we aim to distinguish synchrony that leads only to isolated high-valued entries in the intersection matrix (non-SSEs) from synchrony that leads to DSs. Concretely, the steps are:

I. Generate S (e.g. $S = 1000$) surrogates of the binned discretized data by shuffling the bins randomly.

II. For each surrogate s, $s = 1, 2, \ldots, S$, compute the corresponding intersection matrix $\tilde{I}_s$ and derive the associated probability matrix $\tilde{P}_s$ by Equation (4.2) or Equation (4.3), using the original firing rate profiles. This operation preserves all synchronous events in the original data but not their temporal order, and effectively corresponds to an identical random shuffling of rows and columns of P. Any DS originally present in P is thus destroyed in $\tilde{P}_s$, thereby implementing the null hypothesis $\mathcal{H}_{00}$ that the observed repeated events do not form temporal sequences.

III. Transform each surrogate probability matrix $\tilde{P}_s$ into a filtered probability matrix $\tilde{F}_s$ as follows:

   – apply to each entry $\tilde{P}_{s;ij}$ a rectangular kernel

   – for each $\tilde{P}_{s;ij}$ extract its d largest (most significant) neighbors $\tilde{P}^{(k)}_{s;ij}$, $k = 1, 2, \ldots, d$ and set $\tilde{F}_{s;ij} := 1 - \prod_k \left(1 - \tilde{P}^{(k)}_{s;ij}\right)$.

IV. In the same way transform the original probability matrix P into a filtered matrix F.

V. Compute the significance of each entry F_{ij} in F by comparison with the sample $\{\tilde{F}_{s;ij},\ s = 1, 2, \ldots, S\}$, i.e. set

$$\tilde{J}_{ij} = \frac{1}{S}|\{s :\ \tilde{F}_{s;ij} \leq F_{ij}\}|.$$

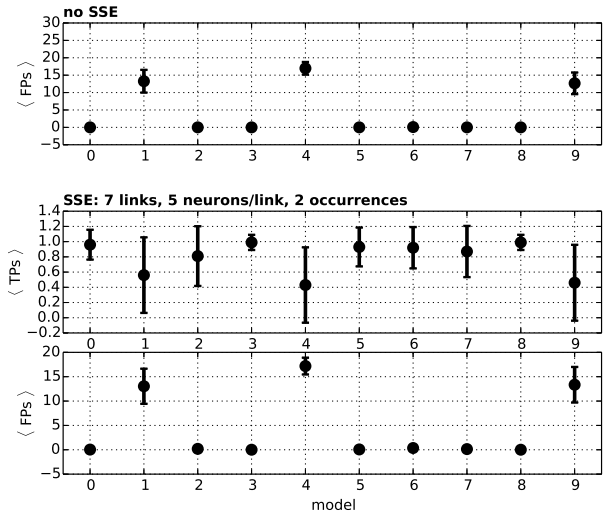

Figure 4.10: **Performance of the Monte-Carlo approach. Top panel:** Average number of FPs (vertical axis) found over 100 simulations of each model from 0 to 9 (horizontal axis) using the Monte-Carlo estimate $\tilde{J}$ of the joint probability matrix under $\mathcal{H}_{00}$. **Middle panel:** Average number of TPs found for each model from 0 to 9 after injecting two repetitions of an SSE (equivalent to one DS in the intersection matrix) with 7 events and 5 neurons/event. **Bottom panel:** Average number of FPs for the same data as in the middle panel.

The matrix $\tilde{J}$ is a Monte-Carlo estimate of the joint probability matrix under $\mathcal{H}_{00}$. Its entries $\tilde{J}_{ij}$ are statistically significant if $\tilde{J}_{ij} > \alpha_2$. In summary, if a repeated SSE was present in the data, the surrogate destroys the corresponding DS in I and thereby leads to higher (more significant) values in $\tilde{J}$ than if there was no such arrangement in the first place, as e.g. for non-SSE synchrony.

Figure 4.10 shows the performance on models 0 to 9 using this Monte-Carlo estimate (and thus working under $\mathcal{H}_{00}$). The FP rate is comparable to that obtained using the analytical approach under H_0. Importantly, this holds also for models $6-8$ where correlations are present, demonstrating that the correlations embedded in the data do not bias the analytical estimates. The Monte Carlo approach shows also decreased TP rate, especially for the cases where firing rates are non stationary (models 1, 4 and 9, TP rate < 0.6). The reason is that time bin shuffling effectively destroys the rate profiles and thus ignores rate changes when building the null distribution. Nevertheless, scenarios may be thought of where correlations need to be taken into account by means of the Monte-Carlo approach illustrated above.

Time-varying firing rates. We further investigate how time-varying firing rates affect the performance of the method. Model 1 contains independent Poisson spike trains, as in model 0, with the difference that the firing rate profiles exhibit a coherent rate jump from 10 Hz to 60 Hz at time $t_1 = 600$ ms and back to 10 Hz at time $t_2 = 700$ ms (see Figure 4.5A and Figure 4.6). Similarly, models 4 and 9 combine this rate non-stationarity with Gamma-distributed ISIs (Figure 4.5C) and with cross-neuron rate heterogeneity (Figure 4.5B), respectively. In all cases the FP rate increases beyond 14.0 (*filled circles* in Figure 4.9). The problem resides in the inability to estimate the instantly varying firing rates by means of a fixed-width kernel convolution, which smooths the rate jump over a larger window. Indeed, when using the true rate profiles to compute the statistics, the FP rate drops back to values lower than 2.0 for models 1 and 9, and as low as 5.7 for model 4 which features highly regular ISIs (*empty circles*).

In realistic scenarios of data analysis of experimental data firing rate profiles are not known. Often though, trials, i.e. repetitions of the same stimulus or behavioral condition related to a given trigger event, are available that make the estimation of the firing rate possible and more reliable. Common tools for firing rate estimation across trials are the

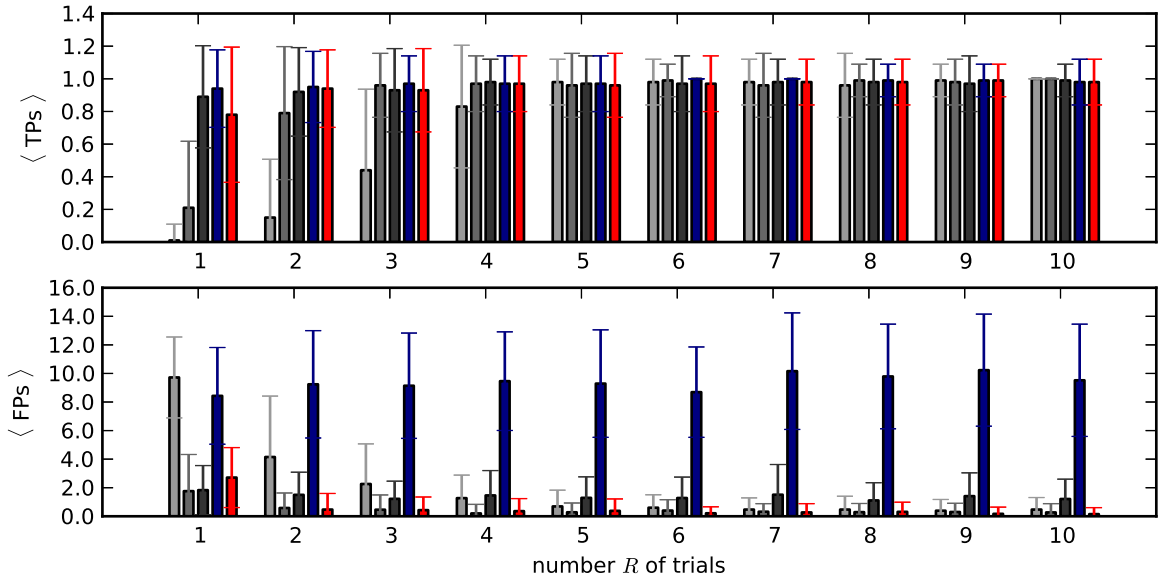

Figure 4.11: **Rate profiles estimated via 3 different methods computed over an increasing number of trials.** Average number of TPs (*top*) and FPs (*bottom*) resulting by different firing rate estimation methods, as a function of the number of trials used to estimate the firing rates of individual neurons (horizontal axis). The data are defined by model 4 and enriched with two repetitions of an SSE with $l_{\mathrm{SSE}} = 7$ links and $\xi_{\mathrm{SSE}} = 5$ neurons/link. *Gray bars*: PSTH (dark: 20 ms bin width, medium-dark: 10 ms, light: 5 ms). *Blue bars*: kernel convolution with fixed kernel width (200 ms). *Red bars*: optimized-bandwidth kernel convolution.

peri-stimulus trial histogram (PSTH, Perkel et al., 1967a) or trial-averaged kernel estimates. We use model 4, which represents a worst case scenario (sudden comodulation of rates, non-Poisson spiking), to investigate whether cross-trial estimation of the firing rates is accurate enough for our needs since it comprises a worst case scenario (coherent rate jump, non-Poisson). To this end we estimate the firing rate profile of each neuron by simulating its activity R times ("trials"), $R = 1$ through 10 and computing its time-resolved average rate over these trials using three different estimation techniques (see Section 4.2.5 for details): PSTH with a bin width of 5, 10 or 20 ms, kernel convolution (Nawrot et al., 1999) with fixed kernel width (200 ms) and kernel convolution with an optimized kernel width (Shimazaki and Shinomoto, 2010). Figure 4.11 illustrates the resulting performance for each method as a function of the number R of trials considered for the estimation. The PSTH (grey bars) performs well for a proper choice of the bin width (here the best being 10 ms, resulting in a TP rate $> .9$ and a FP rate $< .2$ when 3 or more trials are available), but the bin width needs to be predetermined. The same problem affects kernel convolution with a fixed kernel width, which is here chosen too large, resulting in large numbers of TPs. Instead, the optimized-width kernel convolution solves the problem by determining the optimal kernel width in a statistical sense, and yields here a TP rate $> .9$ and a FP rate $< .2$ as soon as 3 or more trials are considered. Thus, this proves to be a suitable method for the rate estimation for experimental data, which are typically characterized by strong and fast rate changes, provided that multiple trials are available.

Number of neurons and size of the DS. We finally investigate how the power of the method is affected by a reduced number of neurons involved in SSE activity in the total

data set. To this end, we repeatedly simulate background activity as in model 0 varying the total number N of neurons from 50 to 500, and add two occurrences of an SSE varying the number l_{SSE} of links (from 3 to 7) and the number ξ_{SSE} of synchronous spikes per link (from 2 to 5). For each value of N, l_{SSE} andξ_{SSE} we generate the data 100 times and analyze each of them with ASSET employing the parameters reported in Table 4.2. Note that the kernel length employed ($l_K = 5$) is suboptimal for values of l_{SSE} higher than 3, as in those cases the kernel centered at the ends of the DS does not cover the full DS and a longer kernel would yield higher performance.

The results are shown in Figure 4.12. The TP rate (column A) increases for a larger size ξ_{SSE} of the synchronous events and a smaller total number N of spike trains, as a larger fraction of the total neurons are involved in the SSE. However, the TP rate does not increase substantially with the number l_{SSE} of synchronous events composing the embedded SSE, because the kernel has fixed length and does not cover the additional events, thus resulting in stationary joint p-values J_{ij} of the second statistical test. For low l_{SSE}, low ξ_{SSE} or large N the TP rate drops considerably to values lower than 0.5. This is partly due to our strict definition of a TP, which does not include found SSEs containing less than 50% of the synchronous events composing the embedded SSE (true positive events). Figure 4.12, column B, shows the average fraction of true positive events composing each found SSE. The FP rate (column C) is close to 0 (and always lower than 0.1; note the different scale of the color map) for all investigated values of l_{SSE}, ξ and N. Importantly, it does not increase with N, which is a relevant feature of the method for applications to large-scale data. For $l_{\mathrm{SSE}} = 7$ and $\xi_{\mathrm{SSE}} \geq 3$ we observe occasionally jumps in the FP rate. Almost all of these FPs consist of SSEs which partially overlap with the embedded one, but not enough to be classified as TPs. As a comparison, Figure 4.12, column D, shows the rate of FP SSEs that are completely disjoint from the true one. This rate is indeed close to 0 for all parameter choices, showing that the increased FP rates shown in column C result from partially true discoveries.

We then investigate the special case of SSEs composed of 1 spike per group only ($\xi_{SSE} = 1$), to test if our method is able to detect spatio-temporal patterns involving no spike synchronization. We inject as before two repetitions of a spatio-temporal pattern, now composed of $l_{\mathrm{SSE}} = 7$ spikes with a time delay of 5 ms, into independent data (model 0, composed of $N = 10$ neurons). Because the previous analysis indicates low performance already for $\xi_{\mathrm{SSE}} = 2$ for the parameters employed so far, we increase the kernel length to $l_K = 7$.

As shown in Figure 4.13, the p-value of individual entries in the intersection matrix corresponding to the real DS and the joint p-value of their $d = 5$ largest neighbors are both too large (weakly significant) to yield sufficient test power. A larger number of repetitions of the spatio-temporal pattern may increase the statistical significance to acceptable levels, as discussed in Section 4.4. Thus, we conclude that ASSET is not suitable for the detection of spatio-temporal patterns involving no spike synchrony.

4.3.3 Analysis of an SFC network simulation

SFC data. Schrader et al. (2008) produced synthetic data from simulations of a balanced neuronal network with embedded overlapping SFCs (SFCs, Abeles, 1991). The full network comprises 40,000 excitatory and 10,000 inhibitory neurons, with biologically realistic connectivity. 50 SFCs are embedded, each obtained by selecting 2000 neurons of the network and randomly organizing them in 20 successive groups of 100 neurons each. All neurons in one group are connected to allneurons in the next group in a feed-forward fashion, resulting in high synaptic convergence and divergence. Individual neurons partic-

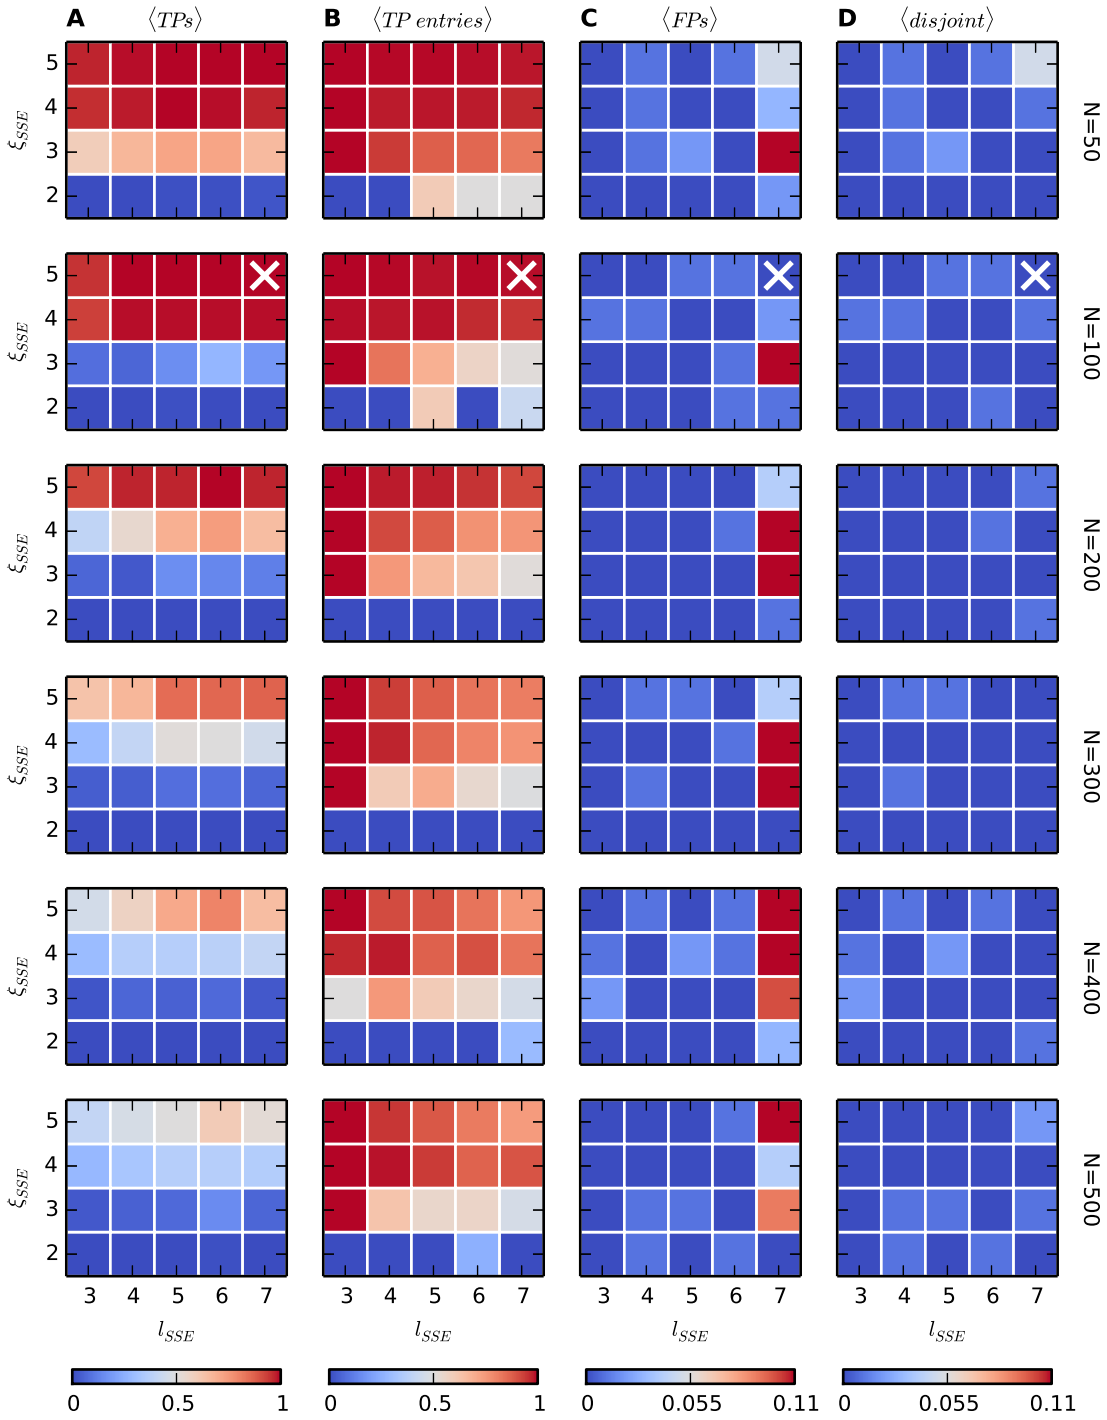

Figure 4.12: ***Performance as a function of*** l_{SSE}, ξ_{SSE} ***and*** N. **(A)** Average number of TP SSEs (over 100 model simulations) found in data where an SSE with varying number of links (*horizontal axis*, from 3 to 7) and neurons per link (*vertical axis*, from 2 to 5) is injected twice in the activity of N otherwise independent spike trains (model 0). N varies from 50 to 500 (top to bottom panels). For $N = 100$ the cross at position ($l_{\mathrm{SSE}} = 7$, $\xi_{\mathrm{SSE}} = 5$) marks the parameters used in the previous calibrations. **(B)** Average number of events of the embedded SSE contained in each found SSE ("true positive events"). **(C)** Average number of FP SSEs (note the different scale of the color map). A FP can either be completely disjoint from the true embedded SSE, or contain part (or all) of its member events (see definition above and Figure 4.8). **(D)** Average number of found SSEs completely disjoint from the embedded one.

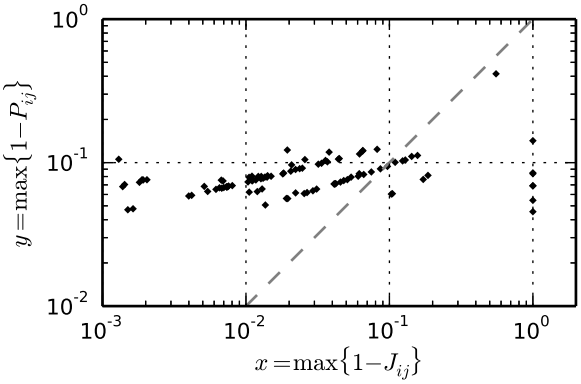

Figure 4.13: **Scatter plot of the p-values of the two tests in the presence of a repeated spatio-temporal pattern.** The p-value y of the first test over the p-value x of the second test for a reduced version of model 0. The model consists of parallel activity from $N = 10$ independent Poisson spike trains, plus two repetitions of a temporal sequence of $l_{\mathrm{SSE}} = 7$ spikes from 7 different neurons ($\xi_{SSE} = 1$). The repeated spatio-temporal pattern generates a DS in the intersection matrix composed of 7 entries $\mathcal{S}_{DS} := \{(i_k, j_k),\, k = 1, 2, \ldots, 7\}$, and corresponding values P_{i_k,j_k} and J_{i_k,j_k} in the probability and joint probability matrices, respectively. The data are generated 100 times. For each data set and corresponding DS in the intersection matrix, the figure shows a point (x, y) given by the maximum (i.e. least significant) of these p-valuesacross all entries composing the DS, i.e. $y = \max_k(1 - P_{i_k,j_k})$ and $x = \max_k(1 - J_{i_k,j_k})$. The figure shows the cloud of the 100 points (x, y), one per simulation.

ipate in up to 3 SFCs, which therefore overlap. The SFCs are stimulated by injection of a current pulse in the first link of the chain, leading to synchronous spiking activity in the first group of the chain which propagates downstream. The propagation delay between one group and the next is ~ 2 ms. The propagation is robust to the presence of noise and does not require all neurons in a link to be active for propagation to the next link (Diesmann et al., 1999).

We consider the activity of 2000 neurons of the network, having identities 8001 to 10,000, over a time segment of 300 ms. The considered neurons compose the entirety of one of the 50 SFCs . In particular, neurons 8001 to 8100 compose the first link, neurons 8101 to 8200 the second link and so on until neurons 9901 to 10,000, which compose the last link of the chain. Figure 4.14 shows the activity of these neurons over a time stretch of 10 s (A) and of 300 ms (B,C). Panels A and B reproduce Figure 2 by Schrader et al. (2008). As apparent from panel B, in the considered time window the selected SFC is active three times. The first and last time the activity runs through all the 20 links, the second time it propagates only until the 13-th link. Panel C reproduces the same data as in panel B, but with a random shuffling of the neuron identities along the vertical axis. Here, the SFC activity misleadingly appears as a co-modulation of firing rates across neurons.

In addition, some of these neurons also belong to 3 successive groups of a second SFC. The 3 groups comprise 14, 12 and 14 neurons of the population analyzed, respectively. In the time considered period, this second SFC is active twice. Contrarily to the first SFC, the activity of the second SFC does not become apparent in the raster plot of Figure 4.14B, because the sorting of neuron identities along the vertical axis of the plot does not place neurons belonging to the same group close-by. Neither does this activity reflect in the plot as population co-modulation, because it involves a relatively small fraction (2%) of the

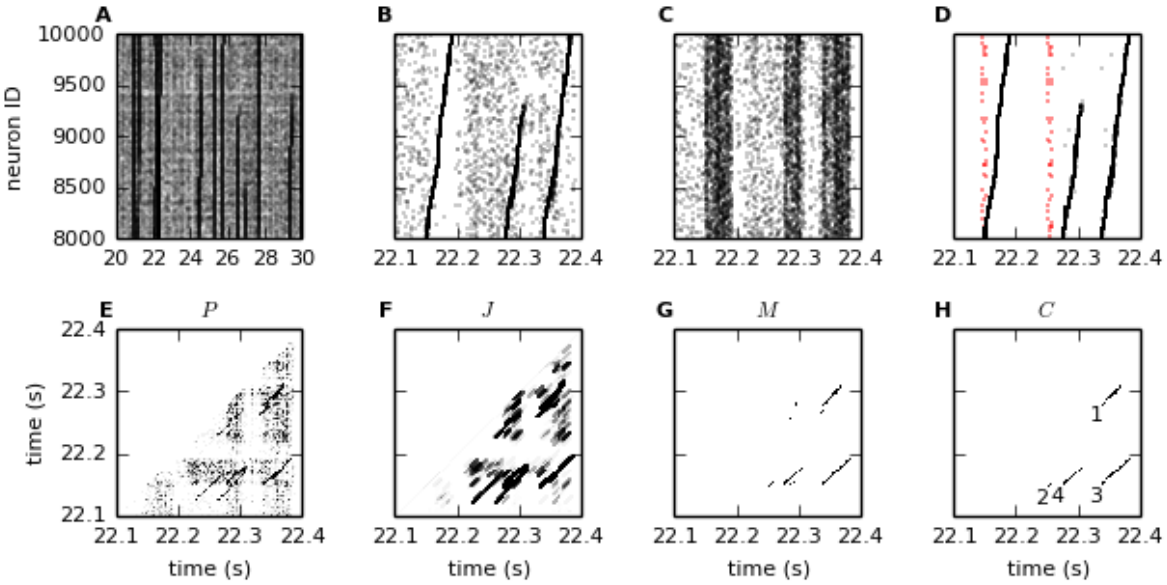

Figure 4.14: **SFC network data and analysis results. (A-D)** Activity of 2000 neurons (numbered 8001 to 10,000) forming the entirety of an active SFC in a balanced random network (Schrader et al., 2008) composed of 20 groups. Neurons $8001 - 8100$ form the first group of the chain, neurons $8101 - 8200$ the second group and so on. A subset of the neurons also participate in 3 successive groups of a second SFC in the network. **(A)** Activity over a time window of 10 s (reproduced from Schrader et al., 2008, Figure 2A). **(B)** Activity over a time window of 300 ms (reproduced from Schrader et al., 2008, Figure 2B). The SFC under consideration is actived three times in this time interval. **(C)** Same as in panel (A), but with random sorting of the neuron ids on the vertical axis. **(D)** SSE activity detected by ASSET. In the time interval considered, the two SFCs are stimulated three times (black dots) and two times (red dots), respectively. **(E-G)** Probability matrix P, joint probability matrix J, masked matrix M and cluster matrix C. Numbers 1 to 3 and 4 in the cluster matrix mark the significant DSs found by ASSET and associated to the activations of the first and second SFCs, respectively.

population.

Analysis of SFC data. We analyze the stretch of data illustrated in Figure 4.14B with ASSET to demonstrate that the method is able to discover repeated SFC activation. The activation of each SFC generates an SSE, and each pair of SSEs corresponds to a diagonal structure (DS) in the intersection matrix. Therefore, the triple activation of the first SFC generates 3 diagonal structures that the method should find in the cluster matrix: one (DS 1, composed of 20 links) is given by the overlap of the two complete runs of the SFC (which involve 20 synchronous events each), and the other two (DSs 2 and 3) are composed of 13 links each, resulting from the comparison of each complete activation with the one stopping at the 13-th group. The double activation of the second SFC, composed of 3 successive synchronous events each, should analogously yield 1 significant DS (DS 4) of length 3. Indeed, the method finds all of these DS, as shown in Figure 4.14H. The raster plot in Figure 4.14D shows the SSEs generated by the activation of the two SFCs (black and red dots, respectively; spikes not belonging to the SSEs not shown), which ASSET reconstructs in their entirety.

Table 4.3, first row, reports for each DS the true number of composing entries and

	DS 1	DS 2	DS 3	DS 4
true nr. entries (dx/dy)	**20**	**13**	**13**	**3**
nr. entries (dx/dy)	23 (14/15)	13 (9/9)	18 (10/10)	3 (2/2)
median i-mat value	93	88	53	14
median p-mat value	$< 10^{-22}$	$< 10^{-22}$	$< 10^{-22}$	$\sim 6.6 \cdot 10^{-8}$
median j-mat value	$< 10^{-22}$	$< 10^{-22}$	$< 10^{-22}$	$\sim 3.9 \cdot 10^{-14}$

Table 4.3: **Diagonal structures of active SFCs.** The top row shows in bold the true number of repeating synchronous events forming the four pairs of repeating SSEs present in the data. The second row shows the number of entries composing each of the four associated DSs as found by the analysis, and in brackets the distance (in number of bins) between the first and the last event of the first SSE occurrence (dx) and the second SSE occurrence (dy). The other rows in the table show the median value of the entries in the intersection matrix (*second row*), in the probability matrix (*third row*) and in the joint-probability matrix (*fourth row*) corresponding to each DS.

the true distance (in number of bins) between the first and the last event for the first SSE occurrence (dx) and for the second SSE occurrence (dy). The second row shows the corresponding values found by the analysis. Due to a larger bin width (3 ms) than the inter-link propagation delay ($\sim$ 2 ms), the found SSEs are shorter than the number of links in the chain: some successive events fall in the same time bin and are therefore merged into a single event. Nevertheless the composing spikes are correctly retrieved, as well as the associated units. Rows 3 to 5 in the table show the median values of matrix entries corresponding to the DS in the intersection matrix, the probability matrix and the joint probability matrix, respectively. As a comparison, the median values of the full matrices were 0, 1 and 1 respectively, indicating the high statistical significance of the entries forming the found DSs compared to the other entries in the matrices. Note that the kernel length employed (length $l_K = 5$) is shorter than the number of entries composing the DSs associated to the first SFC (≥ 9). Nevertheless, these DSs are successfully retrieved.

These results taken together demonstrate the ability of the method to retrieve repeated SSE activity, here generated by repeatedly active SFCs, in massively (here, 2000) parallel spike trains. The method yields high significance values even when when employing sub-optimal parameters, such as larger bin size than the inter-link transmission delay or significantly smaller kernel length than the DS length.

4.4 Discussion

Temporal sequences of synchronous spike events (SSEs) have been postulated as a working mechanism of activity propagation in the cortex (Abeles and Gerstein, 1988; Prut et al., 1998; Ikegaya et al., 2004). The present chapter introduces a novel statistical method for the detection of SSEs in MPST data, named ASSET (Analysis of Synchronous Spike EvenTs). The method is inspired by a visual technique first proposed by Schrader et al. (2008), which represents the repeated occurrence of an SSE as a sequence of large entries along the diagonals of an intersection matrix (diagonal structure, or DS) that indicates for any two time bins the number of neurons firing in both bins. ASSET automatizes the detection process of the original visual technique, assesses the statistical significance of SSEs by exploiting the multiple evidence of its events to derive their joint significance, and determines the structure and neuronal composition of the identified SSEs. In evaluating

the null distribution, the method accounts for the temporal profiles of the firing rates of the observed neurons. As such, it detects SSEs which cannot be explained on the basis of rate coding mechanisms, and thus arise from spike correlations on a shorter temporal scale. Rate correlation is understood as a conceptually different mechanism of information processing than spike synchrony (as for SSEs), because it corresponds to stochastic (probability < 1) rather than reliable (probability 1) neuron activation (De la Rocha et al., 2007; Shea-Brown et al., 2008; Staude et al., 2008; Schultze-Kraft et al., 2013).

We assess the performance of ASSET in terms of false positive (FP) DSs found in various types of stochastic models which mimic typical features of neuronal spike trains, such as variable firing rates in time or across neurons, different inter-spike interval distributions, and correlation structures differing from SSEs. We then additionally inject repeated SSE activity in the data to assess the power of the method in terms of true positive (TP) detections. The analysis requires to perform a statistical test for each pair of time bins, which amounts to tens of thousands of tests for our stochastic simulations of 1 s binned at 5 ms. To avoid FPs, the statistical threshold needs to be set to low values (here, 10^{-5}). This is possible without incurring into large levels of FNs because the joint p-values that the method evaluates are very low (typically $< 10^{-5}$, and as low as 10^{-13}) for the embedded SSEs. Indeed, the method shows high performance, i.e. FP and FN rates both close to 0.

The underlying null hypothesis assumes independent Poisson spike trains, which enables an analytical formulation of the test statistics. The method proves to be robust to deviations from Poissonianity, such as a higher regularity of the inter-spike intervals, which may be observed in experimental data (Shinomoto et al., 1999; Nawrot et al., 2000). It is also selectively sensitive to SSEs, but not to other models of spike correlation, such as synchronous events not organized in a temporal sequence. Nevertheless, anticipating scenarios where strong correlations not forming an SSE might indeed bias the statistics, we propose a Monte-Carlo approach to account for these correlations. To this end, we construct the null hypothesis by estimating the probability to find a certain degree of pattern overlap from the repeated generation of surrogate intersection matrices. In our test data the Monte-Carlo approach yields results comparable the analytical approach, yet at a considerably higher computational cost.

Furthermore, the method is able to distinguish SSEs from repeated precise temporal sequences of sharp, local rate transients from one group of neurons to another (rate propagation). Rate peaks increase the probability of the involved neurons to spike, which remains nevertheless a stochastic event, in contrast to SSE activity. Propagation of rate transients thus generates spike patterns that are different in composition (due to the stochastic activation of neurons) and less precisely timed than SSEs, but more closely resemble the latter as the rate modulation becomes higher and faster. It is likely that, for extremely high firing rates and very short rate jumps, ASSET would not distinguish the two models, although rate coherence and spike synchrony were shown to be mathematically different concepts (Staude et al., 2008).

ASSET critically relies on the estimation of the firing rate profile of each neuron to compute the expected overlap of neuron ids and thus to estimate the p-value of the observed overlap. Estimating firing rate profiles on the basis of single trial spike data typically requires some kind of sliding window approach, such as convolution of each spike with a temporal kernel (Nawrot et al., 1999). Temporal averaging smears out peaks of the underlying original firing rate that occur on a shorter time scale than the window width, and creates artificial peaks if the window width is excessively short. We demonstrate that the single-trial rate estimates obtained by kernel convolution in the presence of time-stationary firing rates yield high performance of ASSET, but lead to impaired performance when the rates are non-stationary in time on a fast time scale, in particular if the rate excursion

is coherent across neurons. The reason is that smoothing by convolution underestimates positive rate peaks and thus the expected overlap, yielding FPs. Estimating the firing rates on the basis of trial averages solves the problem. We here test three such approaches, namely the peri-stimulus time histogram (PSTH, Perkel et al., 1967a), a trial-averaged kernel convolution with fixed kernel width (Nawrot et al., 1999) and a trial-averaged kernel convolution with an optimized kernel width (Shimazaki and Shinomoto, 2010). Our calibration results show that already a small number of trials reduces the FP rate considerably for all three methods, although best performance is found for the optimized-width kernel convolution.

Importantly, cross-trial rate estimation works under the assumption of identical rate profiles across trials. Deviations from this assumption lead to a wrong estimation of the rate in single trials, that is required to calculate the probability matrix, and thereby enhances the FPs. This bias is amplified if the neurons exhibit cross-trial variability in a coherent manner (Grün et al., 2003). Latency variability is a special instance of cross-trial non-stationarity which causes a mis-estimation around the rate onset. In some cases the onset variability of rates can be corrected for by choosing a more proper alignment of trials, e.g. to the stimulus or behavioral event related to the rate change (Grün et al., 2002c). If this is not possible, we suggest to generate the probability matrix P on the basis of surrogate spike data, e.g. by spike train shifting or spike dithering (Grün, 2009). The details of such an approach, however, still need to be explored.

We further investigate how the performance of ASSET relates to various other parameters of the SSEs, such as the number of sequential synchronous events composing the SSE, the number of neurons in each synchronous event and the total number of observed neurons. SSEs are statistically more significant and therefore easier to detect when they involve a larger fraction of the total neurons. However, they can be retrieved even when employing sub-optimal parameters such as a kernel length lower than the length of the DS, given that the SSE involves enough neurons. In contrast, the method does not detect spatio-temporal patterns, i.e. a special case of SSE where each event is composed by a single spike only. Spatio-temporal patterns in a more general sense (with different time intervals between spikes) were suggested as signatures of synfire-chain (SFC) activity in data of low numbers of simultaneous recordings (Nadasdy et al., 1999; Ikegaya et al., 2004; Shmiel et al., 2006).

Bienenstock (1995) proposed a less constrained model of cortical processing (synfire braids) that incorporates synchronous input to individual neurons as in the SFC model, however transferred by connections of different temporal delays compensated by differences in their activation times. Izhikevich (2006) further analyzed this model, which he termed polychronization. Since spike synchrony occurs with a temporal lag in this framework, one expects that, although ASSET is not designed to capture this type of coordinated spiking activity, it may still detect signatures of such activity by choosing a correspondingly larger bin width. We aim to explore such scenario and other applications of ASSET to data from different types of network simulations that exhibit correlations on a fine temporal scale to study the potential of our method in identifying features of the underlying network model.

Importantly, increasing the number of neurons to be analyzed does not increase the computational cost of the method. Indeed, p-values of SSEs are obtained from expressions involving a mere sum of the firing rates of individual neurons, which is virtually instantaneous to evaluate. Rapid advances in electrophysiology will soon enable the simultaneous recording of thousands of neurons (Schwarz et al., 2014). These data promise to expose concerted mechanisms of neuronal coding that remained invisible so far. Our method is designed to keep up with these advances, and to be applicable to the next generation of large-scale recordings of spike data.

If an SSE occurs more than two times, it generates multiple DSs in the intersection matrix, each corresponding to one pair of occurrences. Gerstein et al. (2012) suggested to compute the overlap between triplets (or n-tuples) of bins rather than pairs, and generate a corresponding n-dimensional intersection matrix to visualize DSs in three (or n) dimensions. It is possible to extend this approach to ASSET and to exploit this higher-order evidence to increase the power of the method. This extension will be considered in future work.

Finally, we demonstrate the efficacy of ASSET on data of large-scale simulations of a balanced random network with embedded SFCs, which Schrader et al. (2008) previously generated and analyzed with their original visual technique. ASSET fully reconstructs the SFCs active in the considered time period and does not return additional FPs. Differently from the original technique by Schrader et al. (2008), SSEs are here determined on the basis of their statistical significance and the results are obtained by an automated analysis workflow.

Taken together, these results demonstrate that ASSET is a reliable and effective tool to detect and identify repeated sequences of synchronous spiking activity in MPST data. Whether synchrony propagation is a mechanism of information processing in the neuronal circuitry remains an open question to day, that belongs to the more general debate about the role of fine temporal coding versus rate coding. Convincing theoretical arguments as well as experimental evidence have been provided for both processing mechanisms (see e.g. Roy and Alloway, 2001; Ohiorhenuan et al., 2010; Hong et al., 2012 vs Shadlen and Newsome, 1994; Shadlen and Movshon, 1999; London et al., 2010). However, observing fine-scale temporal correlations requires the simultaneous observation and analysis of sufficiently large portions of the involved neuronal circuitry: the severe subsampling of such circuitry characterizing most data available so far prevents this analysis. Next-generation recording technology promises to expose this concerted activity, whose analysis may finally resolve the long-standing question about the role of millisecond-precise spike correlations in cognitive processing. Our work gives a contribution in this direction by providing, to the best of our knowledge, the first tool for a statistical analysis of synchrony propagation applicable to data of hundreds or thousands of simultaneously recorded neurons.

In future work we plan to apply the method to the reach-to-grasp data described and analyzed for patterns of synchronous spikes with SPADE in Chapter 3. A difficulty that prevents a direct application of ASSET to these data is the variable response latency of the neurons to the stimuli across trials. We compute the population activity within a single trial by a time histogram which counts the total number of spikes across all neurons falling in bins of 20 ms. Figure 4.15 shows heat maps of the population activity for one sub-session over time (horizontal axis) and across trials (vertical axis) for the four trial types. Triggers are indicated by white dots (for details on the experiment and the data, see Section 3.1.1). Panels A and B show the results for trials aligned to the trial start (TS, or equivalently the GO signal provided 2100 ms later), and to the switch release (SR), respectively. In panel A trials are sorted within each trial type by increasing response latencies, quantified by the time delay between the GO signal and the SR. This ordering makes clear that the modulation of population activity around the movement onset is not locked to TS or GO. Similarly, panel B shows that such modulation is not locked to the movement onset either: different trials have different onsets of increase in activity. Because i) estimating firing rates is a required step of the analysis, ii) time-varying rates can be properly estimated via analytical methods only across multiple, stationary trials, and iii) the data under study strongly fluctuate over time and across trials, we need to find another strategy. One possibility would be to split trials into sub-groups with a similar onset of modulation of the population activity. As shown in Section 4.3.2, few stationary trials

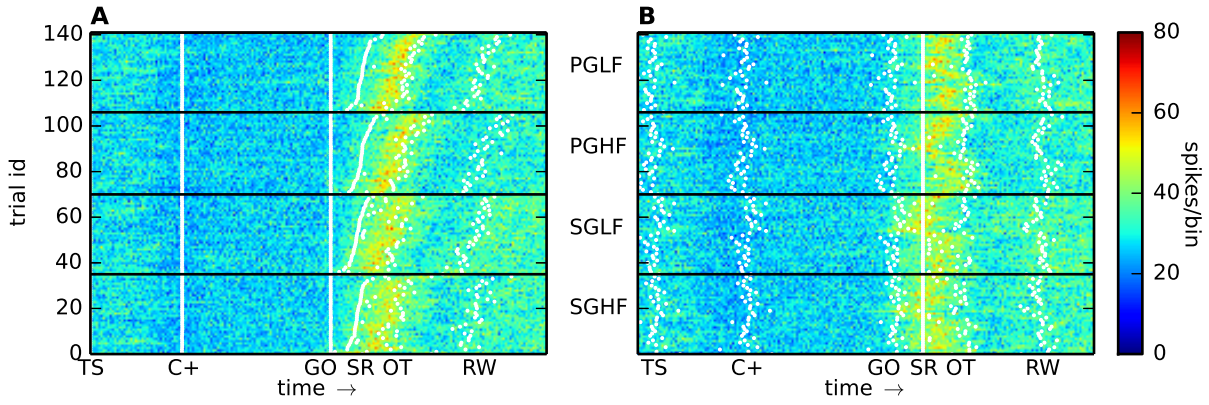

Figure 4.15: **Single-trial population activity in reach-to-grasp data.** Heat maps of population activity in reach-to-grasp data (sub-session i140703-001) over time (horizontal axis) and across trials (vertical axis), separately for each trial type. **(A)** Trials are aligned to the trial start (TS) located at time 0 s, and sorted within each trial type by increasing movement onset, indicated by switch release (SR). **(B)** Trials are aligned to SR. The population activity is computed for each trial by discretizing time in consecutive bins of 20 ms, and by counting the total number of spikes falling in each bin (summed across the 150 neurons recorded during the sub-session). The white dots mark task-related behavioral events (compare with Figure 3.1): beside TS, GO and SR, the visual cue onset (C+), the object touch (OT) and the reward (RW).

would be sufficient to draw a good estimate of the firing rates and consequently to achieve a high performance of the method. A second possibility may be to estimate rates on single trials by means of surrogate methods, as done for the SPADE analysis in Chapter 3. This approach needs to be explored. Once the single-neuron properties of the data will be fully characterized, we will be able to investigate the data with ASSET for the emergence and computational role of SSEs in motor cortex in relation to behavior.

Chapter 5

Conclusions and perspectives

5.1 Questions and answers, in short

The scope of this thesis was to enable the analysis of synchronous spiking activity in massively parallel spike train (MPST) data, where the large number of active neurons and therefore of patterns to consider prevents the use of standard techniques effective on small data.

Spike synchronization at millisecond precision has been shown to take place in relation to a number of behavioral and cognitive processes (for a review, see Chapter 1). However, most studies considered data from small numbers of parallel recordings, and accordingly performed pairwise or low-order correlation analyses of those. In this context we addressed three fundamental, yet unresolved questions:

1. Would data from larger populations of neurons confirm that spike synchronization at millisecond precision plays a role in information processing in the cerebral cortex?

2. If so, how does spike synchronization relate to behavior?

3. If spike synchrony plays a computational role, do temporal sequences of synchronous events represent a mechanism of information transer, as described for instance by the synfire chain model (Abeles, 1991)?

To answer these questions, classical methods for the analysis of spike correlations cannot be employed. Indeed, the application of these methods to high-dimensional data as those provided by modern electrophysiological techniques is impossible, or prone to computational and statistical challenges that make the results unreliable. The reasons behind these limitations lie in the inability of classical methods to cope with the explosion in the number of patterns typical of high-dimensional data, as discussed in Chapter 1 and summarized in the next section. New solutions need to be found.

In Chapter 2 we introduced a novel method, called SPADE, designed to search MPST for patterns of synchronous spikes that repeat more often than one would expect by chance if the neurons were independent. SPADE provides the possibility to address the first two questions raised before. In Chapter 3 we applied the method to data from the motor cortex of two monkeys, collected during the execution of an instructed-delay reach-to-grasp task. The analysis allowed us to determine the emergence of spike synchronization in relation to behavior, and to get new insights on its spatial and temporal arrangement, as well as on its locking to specific behaviors.

In Chapter 4 we further introduced a second method, named ASSET, to analyze propagation of synchronous activity through groups of different neurons. Theoretical studies

suggest that synchrony propagation is a plausible mechanism of information transfer and processing in the neural circuitry, compatible with the biological constraints of the brain. Built on top of a visual technique first proposed by Schrader et al. (2008), ASSET is the first fully statistical method for the analysis of synchrony correlation applicable to MPST data, and thus providing the possibility to answer the the third question posed above.

In this chapter we summarize our work, focusing on the most salient features and benefits of our methods over already existing analysis tools with a similar purpose, and on the findings provided by the analysis of MPST data from populations of over 100 neurons recorded from behaving animals. Based on these results, we propose future research directions aimed at gaining further insights on the relevance of spike coordination at fine temporal scale for information coding in the brain.

5.2 Challenges in the analysis of spike synchronization in MPST data

With rapid advances in electrophysiological technology, the coordinated activity of large populations of neurons recorded simultaneously has become accessible to scientists. In vivo parallel recordings from tens or hundreds of nerve cells are becoming increasingly common, not only from anesthetized, but from awake, behaving animals as well (Buzsaki, 2004). Steady progress in the field is going to provide access to simultaneous data from thousands of single units or beyond (Schwarz et al., 2014).

The increasing size of the data requires new tools for the analysis of coordinated spiking at high temporal precision. Classical methods such as the UE analysis (Grün et al., 2002a,b), or more recent methods based on maximum entropy modelling (Shlens et al., 2006; Schneidman et al., 2006; Tang et al., 2008; Shimazaki et al., 2012), define the time-varying state of the population examined as a binary population vector defined in each short time bin (a few milliseconds) by the joint state of active (1) and inactive (0) neurons in the bin, as illustrated in Figure 1.4A. These methods determine, by different approaches, whether there are states which occur more frequently than expected under pre-defined assumptions (such as that spike trains can be approximately regarded as mutually independent Poisson processes, or that they exhibit pairwise but not higher-order interactions). The underlying assumption is that the presence of an active assembly of synchronous firing cells within a larger population would boost the occurrence count of the state associated to the assembly, making the state statistically significant. For instance, in this framework each single activation of a synchronous cell assembly composed of neurons 1 to 4 and immersed in a population of 6 neurons would correspond to one occurrence of the state vector $(1, 1, 1, 1, 0, 0)$. Only occasionally, chance firing of neurons not belonging to the assembly (e.g. neuron 6) at the same activation time of the assembly would yield different a different state vector (e.g. $(1, 1, 1, 1, 0, 1)$). If the firing rate of the neurons within the population is small compared to the time scale which defines synchrony, the probability that these chance events occur is negligible, and the occurrence count of state vectors can be used to count the number of times that the assembly was active. For instance, a neuron firing at 30 Hz would have only a probability of 0.01 to fire by chance within a time bin of width 3 ms during which an assembly activates.

However, these rare exceptions become the rule as the size of the full population increases as compared to the size of the cell assembly. The activation of a cell assembly of, say, 4 neurons embedded in a population of 100 neurons would seldom yield the state $(1, 1, 1, 1, 0, \ldots, 0)$, because additional chance 1s are to be expected from some of the remaining 96 neurons (see Figure 1.5). If, for instance, each of these additional neurons has

a probability of only 0.01 to fire by chance together with the assembly, the probability that *none* does is $(0.99)^{96} \simeq 0.38$. As a consequence, most of the bins where the assembly is active are to be expected to contain additional chance spikes from non-assembly neurons. Thus, bins containing spikes from not identical but yet overlapping groups of neurons may reflect the activation of a common underlying assembly.

To determine the significance of a putative cell assembly in large data, one therefore needs to consider the number of times that the involved neurons exhibit synchronous firing, irrespective of the activity of the other neurons. This means that, rather than counting the occurrences of a specific state vector containing 1s at the positions corresponding to the assembly neurons and 0s elsewhere, synchrony analyses in MPST data need to sum over the occurrences of all states having 1s in the former positions, and 1s or 0s elsewhere. In MPST data where no prior information is available about specific subgroups of neurons to be analyzed for significance, the query has in principle to be extended to all possible neuronal subgroups, which are 2^N for a population of N neurons. This fact raises severe computational and statistical problems. First, counting the synchronization of all such subgroups is computationally prohibitive for $N = O(100)$ neurons, even on modern machines. Second, and assuming to have solved the computational problem, a statistical issue raises due to the massive multiple testing to be performed in order to assess significance for each found state. This makes classical methods for the analysis of synchronous spike patterns, such as the Unitary Event analysis and Maximum entropy models, inapplicable to high-dimensional data, and calls for new analysis tools. These and other limitations of the maximum entropy approach suggested by (Schneidman et al., 2006) are addressed and investigated in greater detail in Appendix A.

The detection of synchronous spike patterns composed by more than 2 spikes poses an additional problem, intrinsic to the definition of synchrony. Spikes are understood as synchronous if they lie within a certain temporal distance from each other. This allowed temporal window, or jitter, accounts for the expected imprecision in the timing of the biological mechanisms that give raise to action potentials. In other words, two neurons can be considered co-active even if one emits a spike a few milliseconds before or after the other. The most common approach to detect synchronous spike patterns in parallel spike trains consists in discretizing time into short, adjacent time bins of width w reflecting the allowed jitter (e.g. $w = 3\,\text{ms}$), and in defining spikes falling in the same time bin as synchronous. Time binning is employed in most methods for the analysis of spike synchronization, among which the Unitary Event analysis (Grün et al., 2002a,b), NeuroXidence(Pipa et al., 2007, 2008) and approaches based on maximum entropy models (e.g. Schneidman et al., 2006; Shimazaki et al., 2012; Kelly and Kass, 2012). Binning however has a drawback: two spikes may be closer together than w, and yet fall in separate time bins and thus be classified as not synchronous. This problem is aggravated when assessing synchrony between ensembles of more than 2 spikes, as needed in HOC analyses: in this case, all pairs of spikes in the ensemble need to fall in the same time bin for a correct identification of the complete synchronous event.

5.3 Analysis of spike synchrony with SPADE

Gansel and Singer (2012) proposed a method to look for patterns of synchronous spikes (or spatio-temporal spike sequences) in MPST data in a computationally efficient manner. The method reduces the immense number of existing patterns to a small number of possibly interesting ones by a heuristical approach, and then evaluates the statistical significance of the latters via a Monte Carlo approach. The validation performed in their study shows that the method is effective in rejecting chance patterns (FP rate $< 5\%$), however at the

expense of generally mid to low levels of correct detections (TP rate $< 80\%$). Besides, the heuristical reduction of all patterns to good candidates is not accessible to analytical treatment.

Aiming at establishing an analytically solid methodology applicable to MPST data, we introduced in Chapter 2 SPADE (Spike PAttern Detection and Evaluation). SPADE solves both the computational and the statistical problems raised from the large number of spike patterns in the data by combining methodology from frequent item set mining with novel statistics for pattern significance. The number of occurrences of subsets of synchronous spikes (or *item sets*) that can be extracted from a full set in each time bin (or *transaction*) is evaluated in an efficient way by considering only patterns which are *frequent*, i.e. occur at least a minimum number of times (in our case, twice), and *closed*, i.e. repeat more often than all their supersets (see Figure 2.1A-C). Infrequent patterns are considered chance events irrespective of their size, whereas non-closed patterns are trivially explained by the occurrences of some of their supersets, which are therefore the only ones left to be examined.

SPADE determines the number of occurrences of each pattern either by time binning (Torre et al., 2013), or by a sliding window approach in continuous time (Borgelt and Picado-Muiño, 2014). The second approach pools the spikes from all neurons composing the pattern under consideration into a single compound spike train, and positions a window of width w such that its left end lies at the time of the first spike. If all neurons in the pattern have a spike falling within the window, the corresponding set of spikes is classified as synchronous, the pattern's occurrence count is increased by 1, and the window is slided to the first spike time not covered by the original window. Otherwise, the window is moved to the second spike time. Then the procedure is repeated. This procedure ensures that, whenever all neurons composing the pattern emit spikes all within a distance w from each other, the pattern's occurrence count is increased (unless those spikes where already considered for another occurrence of the same pattern). Both the time binning and the sliding window approaches are compatible with the frequent item set mining methodology exploited by SPADE to count pattern occurrences in an efficient manner. Defining synchrony is continuous time is computationally slightly more intensive, but ensures not to miss pattern's occurrences. This results in a significantly higher performance for SPADE when neurons exhibit jittered synchronous firing (see Figure 2.10).

Once closed frequent patterns are identified, their p-value has to be estimated. The number of closed frequent patterns is typically orders of magnitude smaller than the number of all patterns in the data (see Figure 2.1D). Nevertheless, in MPST data of even just a few seconds it is still usually too large to allow for statistical assessment of all individual patterns via direct tests, because these would amount to a massive multiple testing. To avoid this problem, SPADE first tests for the significance of the pattern's signature, i.e. of the pair $\langle z, c \rangle$ of pattern size z and pattern occurrences c. Only patterns with a significant signature, i.e. for which z and c are significantly larger than expected to be observed by chance in the considered population, are kept as potentially significant patterns. Because the number of different signatures is usually orders of magnitude lower than the number of closed frequent patterns (a few hundreds as compared to tens of thousands, in the electrophysiological data we considered in Chapter 3), individual signatures can be tested for significance directly, possibly applying stardard statistical corrections such as false discovery rate (FDR) correction. The p-value of each signature is determined using a Monte-Carlo approach: surrogates of the original data are generated by dithering each spike randomly around its original position. The dithering amplitude is to be chosen larger than the time scale which defines synchrony (few milliseconds), but small enough not to destroy firing rate modulations (usually occurring at the time scale of several tens

to hundreds of milliseconds).

Patterns whose signature is significant are further tested for mutual significance conditioning on the presence of one other. Only conditionally significant patterns are kept.

We demonstrated the validity of SPADE under typical features of experimental neuronal firing rates, such as variability over time and across neurons. Our calibration performed on synthetic control data shows that the method is able to correctly distinguish injected synchronous spike patterns from chance synchrony, yielding both high TP and low FP levels.

5.4 Analysis of synchrony in macaque motor cortex: new insights and open questions

Using SPADE on simultaneous recordings from monkey motor cortex during execution of an instructed-delay reach-to-grasp task, we investigated Chapter 3 the emergence of patterns of synchronous spikes during behavior. The data, collected over 10 recording subsessions from each of two different monkeys, L and N, consisted of trials (typically over 100 per subsession) of 4 different types, differing by the grip the monkeys had to use to grasp an object and by the force needed to pull the object. For each subsession, we additionally split each trial into 6 time windows of 500 ms duration, called *trial epochs*, concatenated epochs of the same trial type together, and separately analyzed the 24 data sets resulting from the combination of each epoch and trial type. Because each trial type and epoch involved a different cognitive process, this approach allowed us to associate the occurrence of significant patterns of synchronous spikes to behavior.

The analysis shows that the precise neuronal composition of significant patterns of synchronous spikes was highly specific to the trial epoch and to the grip modality, while weakly or not at all specific to the force level. In other words, different epochs and grip modalities involved the synchronization of different groups of neurons. These findings are consistent with earlier studies showing that distinct population of neurons are involved in movement preparation and execution in instructed-delay motor tasks (Wise et al., 1983; Riehle and Requin, 1989; Confais et al., 2012; for a review see Riehle, 2005).

Because many found patterns partially overlapped in neuronal composition, we further grouped them via a clustering algorithm into clusters of similar patterns, and analyzed the cluster (instead of pattern) specificity to epochs, grip modality and force. We found that clusters are weakly specific to all of these behavioral contexts, suggesting that similar but not identical patterns represent different cell assemblies, which may partially shared neurons but coded for different behaviors. In conclusion, different behavioral contexts triggered the activation of different cell assemblies of synchronous neurons, but some neurons participated in more than one assembly.

We also investigated the spatial organization of pairs of neurons extracted from each significant pattern, and showed a preferential alignment of the pairs along the medio-lateral axis. This orientation is perpendicular to the dominant orientation reported for travelling waves of beta oscillations in the local field potential (Rubino et al., 2006), or the orientation of the neuronal flow as revealed by analyses of Granger causality of spiking activity (Takahashi et al., 2015). It is however consistent with the dominant alignment of synchronized pairs of neurons in the motor cortex along the central sulcus (Kwan et al., 1987). This medio-lateral alignment may subserve the tight functional coupling between proximal and distal representations of the upper limb during the performance of complex reach to grasp movements. However, further investigations will be required to confirm this interpretation. Furthermore, preliminary data obtained partly in the same animals using

single intracortical microstimulation on single electrodes on the array while recording from all other electrodes have shown that the preferred direction between stimulus and recording was primarily parallel to the central sulcus (Hao et al., 2015).

The cell assemblies we found were mostly of size 2, despite the fact that SPADE is not limited to find pairwise correlations. A few patterns were found having size 3 or higher, with a maximum of 4 in monkey L and 8 in monkey N. This may suggest than higher-order correlations, if existing, are indeed negligible. Many theoretical and experimental studies suggested than pairwise correlations suffice in capturing the almost totality of the correlations in data. Schneidman et al. (2006) used maximum entropy models to estimate the portion of information that each correlation order delivers about the joint distribution of states of a neuronal poluation. On this basis, this and later studies (see e.g. Tang et al., 2008) concluded that average firing rates and pairwise correlations explain the almost totality of such information in populations of few cells. Schneidman et al. (2006) further extrapolated these results up to populations of ~ 200 cells by interpolation, and concluded on this basis that indeed higher-than-pairwise correlations may be neglected in fitting neuronal activity to stochastic models. Roudi et al. (2009) later disproved this statement by showing that the size of the network considered in those studies was below a critical value (determined theoretically) that would allow such extrapolation. On the basis of our findings on the paucity of groups of synchronous neurons of size 3 or larger, obtained on populations of size close to or higher than 100 neurons, it is tempting to conclude that higher-order correlations indeed do not matter.

However, despite the increased size of the network we analyzed as compared to the previous mentioned studies, populations of ~ 100 neurons still represent a severe under-sampling of any network dedicated to implement cognitive functions in primates. A square millimeter of cortex in macaque monkey covers a volume containing about $100,000$ neurons. The $3.6 \times 3.6\,\mathrm{mm}$ area covered by the electrode arrays we employed for our study therefore corresponds to a network of roughly $1,000,000$ neurons. The populations we sampled is 4 orders of magnitude smaller. Even if the full network would be organized in cell assemblies of large size, the probability to sample more than 2 neurons from the same assembly is low. Assuming to sample N neurons from a population of S cells, all being equally likely to be selected, the probability to sample exactly ξ neurons from a cell assembly of size C is

$$p(\xi) = \frac{\binom{C}{\xi}\binom{S-C}{N-\xi}}{\binom{S}{N}}$$

for any $\xi < C, N < S$. Thus,

$$\begin{aligned} p(\xi+1) &= \frac{(C-\xi)(N-\xi)}{(\xi+1)(S+\xi+1-N-C)}\,p(\xi) \\ &\approx \frac{NC}{(\xi+1)S}\,p(\xi) \quad \text{for } \xi \ll C, N \ll S. \end{aligned}$$

For $C \ll S$, as expected for any realistic size of cell assemblies (up to hundreds of neurons) as compared to the full neuronal population, $p(\xi+1)$ is orders of magnitude smaller than $p(\xi)$. Thus, sampling 3 or more neurons from the same cell assembly is very unlikely as compared to the probability of sampling at most 2 of those neurons. For instance, for $S = 1,000,000$ and $C = N = 100$, $p(3) \approx p(2)/300$.

Another reason behind the paucity of higher-order correlations in our data may lay in the fact that we removed all spikes forming perfectly synchronous events at sampling precision $\delta = 1/30\,\mathrm{ms}$ (see Section 3.1.2). As illustrated in Figure 3.2B for one representative sub-session, these *hyper-synchronous* events occurred in our data at significantly higher

frequency than one could expect given the empirical firing rates under the hypothesis of independence – a valid assumption at sub-millisecond temporal resolution. We removed these spikes assuming that they may be artifacts of unknown origin, whose presence would yield large numbers of FPs. Nevertheless, by doing so we likely removed also real spikes and real synchronous events, and probably reduced the size of some of the patterns we still found as significant by removing a portion of the composing spikes.

In light of these considerations – based on simplifying yet realistic assumptions – the fact that we still found synchronous spike patterns of size 3 or higher out of data from 100 neurons is a strong indication that assemblies of synchronous spiking neurons exist, and may indeed be composed of large numbers of neurons. Steady progress in electrophysiological technology is providing simultaneous recordings from ever larger populations of neurons in vivo (Schwarz et al., 2014), exposing this concerted activity more and more. Analysis tools able to keep the pace with such an increase in data dimensionality will be central in investigating these data for cooperative dynamics. SPADE is a promising tool in this respect, thanks to the frequent item set mining and the pattern spectrum filtering approaches that allow the method to milden considerably the computational load required for pattern search and the statistical problem of multiple testing, respectively.

5.5 Analysis of synchrony propagation with ASSET

Volleys of synchronous spikes from different pre-synaptic neurons impinging on a receiving neuron are considered a plausible mechanism of information coding, because effective in triggering a spike response. In Chapter 1 we reviewed a number of theoretical considerations and experimental findings supporting this hypothesis. Among those, a central role is played by the leaky nature of membrane potentials, which makes synchronous PSPs considerably more effective than asynchronous ones. The potential role of spike synchronization in information processing in the brain motivated the development of the SPADE analysis and its application to electrophysiological data, as discussed above.

Synchrony, however, needs to be spatially and temporally organized in order to constitute a possible mechanism of information coding. Abeles (1991) introduced the synfire chain (SFC) model, which assumes that synchronization propagates through groups of neurons in order to transmit and process information. The activation of a SFC triggers a temporal sequence of synchronous events (SSE), each event involving one group of neurons densely connected to the next group in a highly convergent and divergent manner. The SFC model was shown to be consistent with biological constraints of the neuronal circuitry, and to generate activity which is robust to noise (Diesmann et al., 1999).

Despite attempts to find traces of SFC activity in electrophysiological data from membrane potential traces (Ikegaya et al., 2004), direct evidence has been lacking so far. The main reason for this is the lack of analysis methods able to search high-dimensional spike data for SSEs. Schrader et al. (2008) first worked on this topic, by introducing a visual technique that registers the overlap of spiking activity between all pairs of time bins in a matrix, such that two repetitions of an SSE appear in the matrix as a diagonal structure (DS) of high-valued entries. The method by Schrader et al. (2008) has the merit of being applicable to MPST data with minimal computational effort, but the limit of being a purely visual technique not supported by further quantitative analysis (e.g. assessment of statistical significance of putative DSs).

In Chapter 4 we refined this method by deriving an analytical expression for the p-value of DSs in the matrix under the null hypothesis that spike train are Poisson and independent. Our improved method, called ASSET (Analysis of Sequences of Synchronous EvenTs), evaluates the null distribution of values taken by the entries of the matrix by

accounting for the temporal profile of the firing rate of each observed neuron. As such, the method incorporates rate changes over time and across neurons - two typical features of experimental data, in the null hypothesis. The detected SSEs cannot be explained on the basis of rate coding mechanisms, and thus arise from spike correlations on a shorter temporal scale. Rate correlation is understood as a conceptually different mechanism of information processing than spike synchrony (as for SSEs), because it corresponds to stochastic (probability < 1) rather than reliable (probability 1) neuron activation (De la Rocha et al., 2007; Shea-Brown et al., 2008; Staude et al., 2008; Schultze-Kraft et al., 2013).

Extensive calibration on stochastic test data demonstrated the validity of ASSET on a wide range of scenarios. Importantly, deviations from the assumption of Poissonianity, such as a higher regularity of the inter-spike intervals possible in experimental data (Shinomoto et al., 1999; Nawrot et al., 2000), do not impair performance significantly. Also, ASSET is selectively sensitive to SSEs, but not to other models of spike correlation, such as synchronous events not organized in a temporal sequence or propagating rate waves.

A critical feature of ASSET is the stability of its computational cost for an increasing number of parallel spike trains to be analyzed. Indeed, p-values of SSEs are derived from analytical expressions involving a simple sum of the firing rates of individual neurons. Thus, the method promises to be applicable to future big data consisting of larger and larger simultanoeus recordings. In turn, data of this dimensionality will be necessary to sample detectable portions of SFCs. Note however, that the statistical significance of SSEs drops as their size decreases as compared to the full population. In principle, the simultaneous activity of large populations may contain too much noise for an SSE to become distinguishable. Nevertheless, the low joint p-values of entries forming a DS (typically $\sim 10^{-7}$, and as low as 10^{-13}) in the simulated data we analyzed (100 parallel spike trains) suggest that increasing the size of the population by one or two orders of magnitude would not make the p-values drop below significance threshold.

In contrast, the method does not detect spatio-temporal patterns, i.e. a special case of SSE where each event is composed by a single spike only. Spatio-temporal patterns in a more general sense (with different time intervals between spikes) were suggested as signatures of SFC activity in data of low numbers of simultaneous recordings (Nadasdy et al., 1999; Ikegaya et al., 2004; Shmiel et al., 2006).

Bienenstock (1995) proposed a less constrained model of cortical processing, called *synfire braid*, that incorporates synchronous input to individual neurons as in the SFC model, however transferred by connections of different temporal delays compensated by differences in their activation times. Izhikevich (2006) further analyzed this model, which they termed polychronization. Since spike synchrony occurs with a temporal lag in this framework, one expects that, although ASSET is not designed to capture this type of coordinated spiking activity, it may still detect signatures of such activity by choosing a correspondingly larger bin width. We aim to explore such scenario and other applications of ASSET to data from different types of network simulations that exhibit correlations on a fine temporal scale to study the potential of our method in identifying features of the underlying network model.

In future work we plan to apply ASSET on the reach-to-grasp data from monkeys L and N which we analyzed in Chapter 3 with SPADE for patterns of synchronous spikes. ASSET will allow us to determine the existence of temporal sequences of such events in these data.

5.6 Towards spatio-temporal patterns

In this thesis we have introduced two novel statistical methods for the analysis of spike synchrony and synchrony propagation in MPST data. The interest in spike synchrony lies in the observation that synchronous input impinging on a receiving neuron is most effective in triggering a spike response. Thus, neurons may behave like coincidence detectors (Abeles, 1982b), as confirmed by several experimental studies (Roy and Alloway, 2001; Bender et al., 2006; Fino et al., 2010; Hong et al., 2012, e.g.). The synfire chain model (Abeles, 1991) provides a theoretical framework to embed groups of synchronous neurons in networks capable to perform computations. The model is characterized by "chains" of successive cell assemblies (links) feed-forward in a manner that allows volleys of synchronous spikes to propagate from the first to the last link. The model is consistent with biological constraints, and allows for a robust propagation of activity through the networ in the presence of noise (Diesmann et al., 1999). Waddington et al. (2011) have shown that neural networks can self-organize into synfire chain networks though STDP.

However, stable propagation only relies on synchronous input to the receiving neuron. This condition holds if conduction delays between neurons are different, but compensated by different spike emission times such that the emitted spikes arrive to the receiving neuron at the same time. This scheme, suggested in the synfire braids (Bienenstock, 1995) or the polychrony model (Izhikevich, 2006), allows for activity propagation, however reflected by properly timed spatio-temporal patterns (STPs) rather than SSEs.

Considerations about the structure of the neuronal circuitry (e.g. Abeles, 1982a) support the hypothesis that spiking activity may be organized in millisecond precise STPs. Feedback connectivity generates re-entrant activity through chains of neurons (Edelman, 1993). Learning processes may shape this connectivity such to favor the formation of spatio-temporal patterns. Statistical methods exist for the analysis of STPs in data from a small number of parallel spike trains (see e.g. Dayhoff and Gerstein, 1983a; GL Gerstein, 1985; Abeles and Gerstein, 1988; Aertsen et al., 1989). The analysis of small populations of neurons provided experimental evidence of the occurrence of STPs in cortical slices (Beggs and Plenz, 2004) and preparations (Dayhoff and Gerstein, 1983b), as well as in vivo (Arieli et al., 1995; Villa and Abeles, 1990; Abeles et al., 1993; Prut et al., 1998; Villa et al., 1999). The results, however, remain controversial (Baker and Lemon, 2000): most methods employed assume constant firing rates of the neurons, so that the patterns could be mere FPs resulting from sudden rate modulations. Furthermore, the small size of the data considered does not allow to conclude on STP-related mechanisms of information processing for populations of computationally meaningful size.

New analysis tools are therefore needed for the analysis of STPs in MPST data. Although STPs may be seen as SSEs where each event consists of a single spike rather than an ensemble of coincident spikes, ASSET is not able to detect them, because it requires synchrony to boost the statistical significance of each individual event (see Figure 4.13). In principle though, one may think to extend the method is some way in order to capture temporal spike sequences. This task is not straightforward due to the number of possible spatio-temporal patterns in MPST data, which is orders of magnitude larger than the number of SSEs and thus poses severe computational and statistical problems. Data mining techniques, as the one exploited by SPADE to perform an efficient search of synchronous patterns, will likely be needed, with the necessary adaptations, to make the search for spatio-temporal patterns possible.

Appendix A

Ising maximum entropy models of parallel spike trains (and when they fail)

A.1 Motivation

Using an Ising maximum entropy model (iMEM) fitting average rates and pairwise synchronous events in data (iMEM_2), Schneidman et al. (2006) investigated synchrony in the activity from cell cultures in the retina. By this approach, they concluded that triple- or higher-order correlations (HOCs) were negligible. Shlens et al. (2006) reported similar findings on parasol retinal ganglion cells from primates, claiming that "multi-cell firing patterns in the presence of visual stimulation are explained by pairwise, adjacent interactions". By a theoretical approach, Barreiro et al. (2014) suggested that these findings might be motivated by the unimodal shape of input distributions in low-hierarchy networks such as the retina, suggesting that a binomial distribution of inputs might give raise to HOCs in cortical areas. Since the paper by Schneidman et al. (2006), iMEMs have become a widespread tool to characterize spike synchronization in parallel spike train data, not only in retina (Tkacik et al., 2006), but in cortical regions as well (Tang et al., 2008). Most of these studies advocate for the absence or negligibility of higher-order interactions, concluding that correlations in these data, if present, are explained by 2-ways interactions only. iMEMs have later on been extended in the field to account for time-varying correlations of order 2 or 3 as well as for the previous spiking history of the neurons under study (Shimazaki et al., 2012; Kass et al., 2011; Kelly and Kass, 2012).

Schneidman et al. (2006) concluded on the irrelevance of higher-order correlations on the basis of an index introduced by Schneidman et al. (2003), and here called $\mathcal{R}_\xi$, to quantify the portion of total information in data explained by dependencies of a certain order $\xi \geq 2$. Using the index $\mathcal{R}_2$, they found that pairs of synchronous spikes account for 99% or more of the total information in the data they examined. However, these results were not supported by calibration on ground-truth data, a fundamental missing piece to demonstrate the ability of the employed model to fit the data, as well as the ability of the index $\mathcal{R}_2$ to quantify the importance of higher-order versus pairwise correlations. A similar approach was adopted in later studies (e.g. Tang et al., 2008), still without calibration. Because these models fit time averages of firing rates and pairwise correlations, modulations of these quantities over time may bias the results (Grün, 2009). Beside, even in the absence of bias the sensitivity of the suggested information index to sparse but strong higher-order synchronous events remains to be demonstrated.

In this chapter we replicate the analysis suggested by Schneidman et al. (2006) on parallel spike train data which are independent, or pairwise-correlated, or which contain higher-order synchronous events. The results are instructive:

- One can construct data which contain significant synchronous events of high order, but such that the index $\mathcal{R}_2$ takes values higher than 99%. The reason is that strong but sparse higher-order correlations might have a mild effect on $\mathcal{R}_2$ if pairwise correlations are ubiquitous, thus carrying low information about the pattern joint probability distribution. This demonstrates that the two concepts - significance and "information content" of HOCs - should be kept distinct and quantifying one does not allow to draw conclusions on the other.

- iMEMs of order 2 tend to overestimate pairwise correlations in purely pairwise-correlated data if the true correlations are weak, and analogously to find false positive correlations in independent data.

Criticism was already moved to maximum entropy models in the past years. Roudi et al. (2009) showed that maximum entropy models are blind to HOCs when applied to small subnetworks sampled from larger networks with HOCs. Our approach is different: we show that, even when the full data (of modest size) are considered, iMEMs can fail in detecting existing HOCs.

Section A.2 introduces Shannon entropy and MEMs. Section A.3 presents the data types we employ to calibrate the method. In Section A.4 we fit maximum entropy models of order 1 (independent) and 2 (with pairwise correlations) to various data types, showing the inefficacy of the index $\mathcal{R}_2$ to express the relevance of HOCs even in the presence of good fits. In Section A.5 we show that iMEM_2 may be insensitive to weak pairwise correlations, and suggest some care before blindly fitting such model to data. I leave most mathematical and numerical details to dedicated Appendices.

A.2 Maximum entropy models

A.2.1 Shannon entropy and maximum entropy distributions

Given a finite sample space $\Omega = \{\omega_1, \ldots, \omega_n\}$ and a probability $P : \Omega \to [0, 1] : \omega_i \to p_i$, we write the Shannon entropy (or simply *entropy*) of P as

$$\mathcal{H}(P) = -\sum_{i=1}^{m} p_i \log p_i, \tag{A.1}$$

where $0 \log(0) \equiv 0$[1]. $\mathcal{H}(P)$ quantifies the prior uncertainty about the outcome of stochastic sampling from Ω via the probability distribution P. The higher is $\mathcal{H}(P)$, the more uniform and, in this sense, uncertain the outcome of the sampling will be. $\mathcal{H}(\cdot)$ has a minimum at 0 when P is concentrated at one point ($p_i = 1$, $p_j = 0$ $\forall i \neq j$, i.e. deterministic sampling) and a maximum at $\log(m)$ when P is the uniform distribution ($p_i = 1/n$ $\forall i$). Shannon (1948) proved that $\mathcal{H}(P)$ is the only positive, continuous function of P which increases as a function of n when P is the uniform distribution.

We just stated an essential fact: among all possible probability distributions over the probability space Ω, the one maximizing Shannon entropy (or the prior uncertainty about

[1] A more general definition of entropy is $\mathcal{H}(P) = -K \sum_i p_i \log_b p_i$, where $K > 0$ and $b > 1$. K is conventionally set to 1, while the logarithm base b is usually set to 2, e or 10. We choose $b = 10$ (this choice doesn't alter our results qualitatively).

the outcome of a stochastic sampling) is the uniform distribution. This formalizes Laplace's *principle of insufficient reason,* roughly stating that *two events are to be assigned the same probability if there is no reason to think otherwise* (Jaynes, 1957). This concept can be generalized, looking for the distribution with maximal entropy *among those satisfying certain constraints*. In this framework, uniform distributions are unconstrained maximum entropy distribution.

Given k functions (or *features*) $\{f_j\}_{j=1}^{k}$ defined on Ω, we define k constraints for probability distributions P on Ω by requiring that

$$\langle f_j(\omega)\rangle_P = \mu_j \ \forall j, \tag{A.2}$$

where $\langle\cdot\rangle_P$ is the expectation determined by P: $\langle f\rangle_P = \sum_i p_i \cdot f(\omega_i)$. In other words we select, among all probability distributions on Ω, only those that yield average values μ_j for the transformations f_j of the sample space. The constraints are said consistent if there exists at least one probability distribution P over Ω satisfying Equation (A.2). Defined the vector function $\boldsymbol{f}(\cdot) = (f_1(\cdot), \ldots, f_k(\cdot))$ and the vector of constraints $\boldsymbol{\mu} = (\mu_1, \mu_2 \ldots, \mu_k)$, it can be proved that for consistent constraints there is one and only one probability P which maximizes Equation (A.1), and that is the Gibbs (or Boltzmann) distribution (Jaynes, 1957)

$$P(\omega) = \frac{1}{Z}\exp\{-\sum_{j=1}^{k}\theta_j \cdot f_j(\omega)\} = \frac{1}{Z}\exp\{-\boldsymbol{\theta}\cdot\boldsymbol{f}(\omega)\}, \tag{A.3}$$

where $\boldsymbol{\theta} = (\theta_1, \theta_2, \ldots, \theta_k)$ is a vector of real-valued parameters and $Z = Z(\boldsymbol{\theta})$ is a normalizing constant (*partition function*):

$$Z(\boldsymbol{\theta}) = \sum_i \exp\{-(\boldsymbol{\theta}\cdot\boldsymbol{f}(\omega_i))\}.$$

The parameters θ_j are uniquely determined by the constraints via the relation:

$$\mu_j = \frac{-\partial}{\partial\theta_j}\log Z(\theta). \tag{A.4}$$

Although the probability distribution defined by Equation (A.3) is well defined, the relation between the parameters $\boldsymbol{\theta}$ and the constraints $\boldsymbol{\mu}$ is usually not available in closed form. Approximation techniques exist to determine the parameters given the constraints (see Section A.2.5).

The quantity $\boldsymbol{\theta}\cdot\boldsymbol{f}(\omega)$ is called the *energy* of ω. States with higher energy are less likely to occur with respect to states with lower energy.

The interest in maximum entropy distributions lies in the fact that, by definition of entropy, they maximize the uncertainty about the outcome of stochastic sampling from the sampling space, given some constraints about the outcome itself. Any other distribution satisfying those constraint would therefore be more specific about the outcome, i.e. would (maybe implicitly) provide when used for sampling - and thus *assume* when used for modeling - additional information about the sample. Paraphrasing Jaynes (1957), *the maximum entropy distribution [...] is maximally non-committal with respect to missing information.*

A.2.2 Ising maximum entropy models

Fitting a maximum entropy model subject to specific constraints to data and quantifying its goodness-of-fit allows one to assess whether the selected features fully explain the

observations or not. In applications to neuronal modeling, features known to play a role in information processing are the firing rates of individual neurons and spike correlations. In these settings, a neuron i is considered as a stochastic electric device whose state σ_i is observed at discrete times. At each time l, the neuron can be active ($\sigma_{i,l} = +1$) or silent ($\sigma_{i,l} = -1$), like a spin in the classical Ising model. We consider a model network composed of N neurons (spins). Thus, the network's states are stochastic vectors $\boldsymbol{\omega} = (\omega_1, \omega_2, \ldots, \omega_N)$ from the sample space $\Omega = \{-1, +1\}^N$, i.e. binary vectors of length N with binary stochastic elements.

The model proposed by Schneidman et al. (2006) constrains the first moments $\langle \boldsymbol{\omega} \rangle := (\langle \omega_1 \rangle, \langle \omega_2 \rangle, \ldots, \langle \omega_N \rangle)$ of the spins to $\boldsymbol{\mu} = (\mu_1, \mu_2, \ldots, \mu_N) \in (-1, 1)^N$ and the second cross-moments $\boldsymbol{\omega} \wedge \boldsymbol{\omega} = (\langle \omega_i \omega_j \rangle)_{i,j}$ to $\boldsymbol{\Sigma} = (\Sigma_{ij}) \in [-1, 1]^{N \times N}$ (with $\Sigma_{ii} = 1$ for consistency), for a total of $N + \binom{N}{2}$ constraints. Substitution in Equations (A.2)-(A.4) yields the maximum entropy probability

$$P^{(2)}(\boldsymbol{\omega}) = \frac{1}{Z} \exp\{\boldsymbol{h} \cdot \boldsymbol{\omega} + \boldsymbol{\omega} \cdot \boldsymbol{J} \cdot \boldsymbol{\omega}^T\}, \tag{A.5}$$

where $Z = \sum_{\boldsymbol{\omega}} \exp\{\boldsymbol{h} \cdot \boldsymbol{\omega} + \boldsymbol{\omega} \cdot \boldsymbol{J} \cdot \boldsymbol{\omega}^T\}$, subject to the constraints

$$\begin{cases} \mu_i = -\frac{\partial}{\partial h_i} \log Z(\boldsymbol{h}, \boldsymbol{J}) & \forall i = 1, \ldots, N \\ \Sigma_{ij} = -\frac{\partial}{\partial J_{ij}} \log Z(\boldsymbol{h}, \boldsymbol{J}) & \forall i, j = 1, ..., N \end{cases}. \tag{A.6}$$

The values $\boldsymbol{\mu}$ and $\boldsymbol{\Sigma}$ of the constraints are inferred from the data: given observed states $\{\boldsymbol{o_1}, \boldsymbol{o_2}, \ldots, \boldsymbol{o_n}\}$ of a neural network observed over n discrete times,

$$\begin{cases} \boldsymbol{\mu} = \langle \boldsymbol{o}_i \rangle_i & , \\ \boldsymbol{\Sigma} = \langle \boldsymbol{o}_i \wedge \boldsymbol{o}_i \rangle_i & . \end{cases}$$

We refer to the model characterized by Equation (A.6) as the Ising maximum entropy model of order 2, iMEM_2.

Fixing the first moments only yields, by analogous arguments, the Ising maximum entropy model of order 1, iMEM_1,

$$P^{(1)}(\boldsymbol{\omega}) = \frac{1}{Z} \exp\{-\boldsymbol{h} \cdot \boldsymbol{\omega}\}, \tag{A.7}$$

constrained by Equation (A.6), first line. iMEM_1 makes the states of different neurons mutually independent. For this reason we also refer to iMEM_1 as the independent model.

A.2.3 Correlation information

Model iMEM_2 is defined by constraining by the first and second moments of all spins and spin pairs. In Section A.2.4 we show that this is equivalent to constraining the neuron's rates and zero-lag correlations. One could analogously constrain correlations up to a higher order, building a hierarchy of maximum entropy models iMEM_1, iMEM_2, iMEM_3, ..., up to the maximum entropy model constraining all correlations in the data, iMEM_N. We already pointed out that iMEM_1correspond to the independent model reproducing the average firing rates. iMEM_N conversely keeps the average correlations of all orders, and is thus defined by the full joint distribution of the data. The probability distributions of the intermediate models iMEM_i converge to the full joint distribution of the data as $i \uparrow N$.

Setting more and more constraints on the correlations limits further and further the class of models among which to look for the maximum entropy one. Thus, the more constraints are added, the lower is the entropy of the maximum entropy model which

satisfies them (and them only). The entropy $\mathcal{H}_i$ of each model iMEM$_i$ therefore decreases as i increases: $\mathcal{H}_1 \geq \mathcal{H}_2 \geq \ldots \geq \mathcal{H}_N$. Equivalently, the entropy difference or multi-information $\mathcal{I}_{(i)} := \mathcal{H}_1 - \mathcal{H}_i$ grows from 0 up to $\mathcal{I}_{(N)} = \mathcal{H}_1 - \mathcal{H}_N$ as i grows from 1 to N. $\mathcal{I}_{(N)}$ represents the total information conveyed by the data, in terms of the decrease in entropy (uncertainty) determined by considering all correlation orders instead of the rates only. The ratio

$$\mathcal{R}_i = \mathcal{I}_{(i)}/\mathcal{I}_{(N)} \tag{A.8}$$

characterizes the fraction of such information conveyed by correlations of order i or lower. Notice that computing $\mathcal{R}_i$ only requires to characterize the maximum entropy models of order 1 and i. Indeed, these two provide the entropies $\mathcal{H}_1$ and $\mathcal{H}_i$ by Equation (A.1), which allows to compute $\mathcal{I}_{(i)}$. The entropy $\mathcal{H}_{(N)}$ of the data (and thus the multi-information $\mathcal{I}_{(N)}$) can be computed directly from the sample pattern probability, without the need to fit iMEM$_N$.

Schneidman et al. (2003) proposed to use the ratio $\mathcal{R}_2$ to assess whether pairwise correlations provide an effective description of the system. A value close to 1 would imply that pairwise correlations successfully capture the correlation structure of the network, whereas higher-order correlations do not convey, or convey only very little, information about the system (i.e., about the system's stochastic properties). Schneidman et al. (2006), Shlens et al. (2006) and Tang et al. (2008) applied this information measure to different network, in the retina as well as in the cortex, reporting values for $\mathcal{R}_2$ ranging from about 0.85 to over 0.9. They concluded that HOCs were negligible in the networks they examined.

The fact that "HOCs provide negligible information" means that the amount of information conveyed by HOCs about the multi-variate distribution of the population, quantified by $1-\mathcal{R}_2$, is negligible compared to that provided by pairwise correlations and firing rates. It is tempting to summarize this statement by saying that "HOCs are negligible", in the sense that "the data do not contain statistically significant HOCs". This formulation is misleading, and the two facts should be kept separate: HOCs may be present and highly-significant, but nevertheless such that they do not affect the information conveyed by the activity of the full population. In Section A.4 we provide examples of simulated data with an index $\mathcal{R}_2$ close to 1, but containing highly-significant HOCs engaging a large portion of the full network.

A.2.4 Relations between various statistics in Ising maximum entropy models

An Ising spin takes values -1 and $+1$. Therefore, unlike the $0-1$ binary processes we are used to work with, the activity rate (or "firing rate"; the rate of $+1$'s across observations) does not match the spin's mean value (or first moment). Importantly, the spin's mean ranges within $[-1,+1]$ while the rate is constrained within $[0,+1]$. Analogously, the second moments and the correlation coefficients differ.

Given the rate r_i of a spin ω_i, the probability for ω_i to take values $+1$ and -1 are r_i and $1-r_i$ respectively. The relation with the spin's mean $\mu_i := \mathbb{E}[\omega_i]$ readily follows:

$$r_i = \frac{1+\mu_i}{2} \tag{A.9}$$

Noting that the variance of ω_i is $\mathbb{V}[\omega_i] := \mathbb{E}[\omega_i^2] - \mathbb{E}^2[\omega_i] = 1-\mu_i^2$, the correlation coefficient C_{ij} between spins two spins ω_i, ω_j with second moment $\mathbb{E}[\omega_i \cdot \omega_j] = \Sigma_{ij}$ is given by:

$$C_{ij} = \frac{\Sigma_{ij} - \mu_i \cdot \mu_j}{\sqrt{(1-\mu_i^2)\cdot(1-\mu_j^2)}}. \tag{A.10}$$

In a Ising maximum entropy model, the model parameters $\boldsymbol{\theta}$ completely determine the occurrence probabilities of all 2^N spike patterns (network's states) (see Equation (A.3), and Equations (A.5)-(A.7) for the special cases of iMEM_1 and iMEM_2). These determine in turn the spins' moments:

$$\begin{array}{ll} \mu_i & = \sum_{\boldsymbol{\omega}} \boldsymbol{\omega_i} P(\boldsymbol{\omega}), \\ \Sigma_{ij} & = \sum_{\boldsymbol{\omega}} \boldsymbol{\omega}_i \boldsymbol{\omega}_j P(\boldsymbol{\omega}) \end{array} \tag{A.11}$$

and thus the firing rates and correlations, expressed by Equations (A.9) and (A.10). Notice that these relations cannot be inverted: the parameters of the model cannot be expressed as a function of the first and second moments (or equivalently of the rates and correlation coefficients) in closed form. However, approximation techniques exists, such as the one described in Section A.2.5 and employed in this work and in the literature we have referenced to.

A.2.5 Constraining maximum entropy models: generalized iterative scaling

Solving analytically the system of equations (A.4) to determine the parameters of a maximum entropy model given its constraints is in general not possible, or extremely difficult. Numerical techniques have been proposed to find approximate solutions for the parameters, such as generalized iterative scaling (GIS; see Darroch and Ratcliff, 1972). Given desired constraints, the basic idea of GIS is to set the parameters to an initial guess, and then to correct the values by an iterative procedure. At each iteration, the model's probability distribution is given by Equation (A.3), and the features to be constrained (given by the left-hand side of Equation (A.2)) can be derived explicitely. By comparison those with the desired values (right-hand side of Equation (A.2)), the parameters are corrected to reduce the error.

We focus here on the model iMEM_2 (or iMEM_1 for the case that $J_{ij} = 0$). Let

- $\boldsymbol{\theta} = (\theta_l)_l$ be the vector of (unknown) $N + \binom{N}{2}$ model parameters $(\boldsymbol{h}, \boldsymbol{J})$ to be approximated,
- $\boldsymbol{\theta}^{(k)}$ be its approximation at iteration k,
- $\hat{\boldsymbol{f}} = (\hat{f}_l)_l$ be the vector of desired feature statistics (which are sampled from the data: $\hat{f}_l = \langle \omega_i \rangle$ if $\theta_l = h_i$ and $\hat{f}_l = \langle \omega_i \omega_j \rangle$ if $\theta_l = J_{ij}$)
- $\hat{\boldsymbol{f}}^{(k)}$ be the vector of feature statistics derived from the approximated model at iteration k.

Let also indicate the correction to each parameter $\theta_l^{(k)}$ by $\theta_l^{(k+1)} - \theta_l^{(k)} =: \Delta\theta_l^{(k)}$. Tkacik et al. (2006) proposed the following rule (also used in Schneidman et al., 2006) to update the parameters:

$$\begin{aligned} \theta_l^{(0)} &= \hat{f}_l, \\ \Delta\theta_l^{(k)} &= -\frac{1}{k+1}(\hat{f}_i^{(k)} - \hat{f}_i) + \alpha \Delta\theta_l^{(k-1)}, \qquad k \geq 1. \end{aligned} \tag{A.12}$$

which has a fixed point for $\hat{f}_i^{(k)} - \hat{f}_i = 0$. The last term keeps memory of the previous correction in the next one, by an amount determined by the parameter $\alpha \geq 0$. We set α to 0.1.

A.3 Data generation

Our goal is to test the goodness-of-fit of the Ising maximum entropy models of order 1 (iMEM_1, accounting for average firing rates only) and 2 (iMEM_2, additionally fitting the observed second moments) to A) independent data, B) data with pairwise correlations only, C) data with pairwise and higher order correlations. Here we explain how we generate these various data types. Unless otherwise stated, in the reminder we work with samples of dimension (i.e. number of neurons) $N = 10$ and size (i.e. number of observations) $n = 10^6$. Each observation is a state ω_k of the Ising model, represented by a vector of dimension N whose elements take values $+1$ or -1. The sample space is $\Omega = \{-1, +1\}^N$.

Independent data. A way to generate independent samples ω_k given the spin rates r_i (or the means μ_i) is to draw each state $\omega_{k,i}$ of ω_k as a Bernoulli random variable

$$\omega_{k,i} = \begin{cases} +1 & \text{with probability } r_i, \\ -1 & \text{with probability } 1 - r_i, \end{cases}$$

independently for each k and each i. An equivalent approach is to generate spike trains as independent Poisson processes and then bin them into sequences of $+1$s (spike present into the bin) and -1s (no spike). Given the desired rates r_i of $+1$s for the binned sequences, the bin size Δt and the desired sample size n, the corresponding spike trains have stationary firing rates $\hat{r}_i = r_i/\Delta t$ over an observation time $T = n \cdot \Delta t$. We take this approach and set $r_i = 0.05$, $\Delta t = 10\,\text{ms}$ and $n = 10^6$, amounting to $\hat{r}_i = 5\,\text{Hz}$ and $T = 10,000\,\text{s}$.

Data with pairwise correlations only. To generate purely pairwise-correlated data, we draw samples from an iMEM_2. By definition of maximum entropy, this gives us full control over the rates and pairwise correlations and ensures that no other assumptions over the data statistics are being implicitly assumed (see Section A.2).

We generate 10 parallel spike trains having firing rates $r_i = 0.05\ \forall i = 1, \ldots, N$ and correlation coefficients $C_{ij} = 0.05\ \forall i \neq j$, by means of an iMEM_2 having moments $\mu_i = -0.9$ and $S_{ij} \simeq 0.82\,\forall i \neq j$ ($S_{ii} = 1\,\forall i$) (see Section A.2.4) and parameters, approximated via GIS (see Section A.2.5 for details), $h_i \simeq -0.51$ and $J_{ij} =\simeq 0.124\ \forall i \neq j$ ($J_{ii} = 0$). Substituting these parameters in the Gibbs distribution given by Equation (A.3), we randomly sample $n = 10^6$ independent Ising states from the space of all possible $2^N = 1024$ states. The effective rates and correlation coefficients of the simulated data are close to the prescribed one, as shown in Figure A.2A-B.

Data with pairwise and higher-order correlations. To generate data containing both pairwise and higher order correlations, we first generate purely pairwise-correlated data by drawing n_1 samples from an iMEM_2 model of $N = 10$ spins, analogously to what described above. Then, we simulate n_2 samples from a SIP process of order $M \leq N$ as

follows: in each sample ω_k, we set the values of the first $M \leq N$ spins to 1 (active), and the value of the other spins $\omega_{k,i}$, $i = M+1, \ldots, N$, to

$$\omega_{k,i} = \begin{cases} +1 & \text{with probability } r_i = 0.05, \\ -1 & \text{with probability } 1 - r_i, \end{cases}$$

independently for each k and each i.

Finally, we pool the two samples. We set the total sample size $n = n_1 + n_2 = 10^6$ as for the independent and pairwise-correlated data. The rate of the neurons that are not involved in SIP activity is $r_i = 0.95$. The rate of the SIP neurons is given by $(r_i \cdot n_1 + 1 \cdot n_2)/n > r_i$. The pairwise correlations are slightly altered as well with respect to the purely pairwise-correlated data: the SIP neurons have average pairwise correlations $(C_{ij} \cdot n_1 + 1 \cdot n_2)/n > C_{ij}$, because they are perfectly correlated during the SIP epoch, while the non-SIP neurons have average pairwise correlations $C_{ij} \cdot n_1/n < C_{ij}$ because during the SIP epoch they are modeled as independent Bernoulli processes. These differences are negligible if $n_2 \ll n_1$. However, n_2 has to be sufficiently large to make HOCs significant. We set $n_1 = 0.999 \cdot n$ and $n_2 = 0.001 \cdot n = 10000$, which increases the rate of SIP neurons with respect to the pairwise model by about 2% to ~ 0.051. This makes also 2% of the spikes of SIP neurons all-to-all synchronous, which in a population of 10 neurons corresponds to a highly significant event.

A.4 Maximum entropy models fit to data: some tests

A.4.1 Goal and approach

Here we test the goodness-of-fit of models iMEM_1 and iMEM_2 to A) independent data (where iMEM_1 should fit well and iMEM_2 should collapse to iMEM_1, i.e. with coefficients $J_{ij} \equiv 0$), B) data containing pairwise correlations only (where iMEM_1should perform poorly and iMEM_2 optimally) and C) data containing pairwise and higher-order correlations, where neither of the two models should yield an optimal fit and where the coefficient $\mathcal{R}_2$ should be lower than 1 to prove a valid measure of goodness-of-fit for pairwise versus higher-order models.

The procedure used to fit the data is as follows: we estimate from the data the sample statistics (means μ_i and second moments S_{ij}) and apply GIS (see Section A.2.5) to fit the two iMEMs. For iMEM_1 we fit only the means, for iMEM_2 we fit the second moments as well. For each model and data set, we iterate the scaling procedure $L = 1000$ times (some tests with $L = 10000$ yielded very similar results, yet at a $10\times$larger computational cost). We quantify the overall goodness-of-fit of iMEM_ξ by the Kullback-Leibler divergence between the model's probabilities $P^{\text{iMEM}_\xi}(\boldsymbol{\omega})$ of network's states $\boldsymbol{\omega}$ and the observed frequencies $P^{data}(\boldsymbol{\omega})$,

$$D_{KL}(P^{data}||P^{\text{iMEM}_i}) = \sum_{\omega \in \Omega} P^{data}(\omega) \log_{10} \frac{P^{data}(\boldsymbol{\omega})}{P^{\text{iMEM}_i}(\boldsymbol{\omega})},$$

where $0 \log 0 \equiv 0$.

Finally, we quantify the fraction of information conveyed by pairwise correlations through the information index $\mathcal{R}_2$ defined by Equation (A.8).

The results of the fits to the various data types are shown in Figures A.1-A.3. In each figure, the rates and correlation matrix estimated from the data are represented in the small panels on the top right. The results obtained by fitting iMEM_1 and iMEM_2 are shown in the top and bottom row, respectively. From left to right: the ratio between the estimated model means $\mu_i^{iMEM_i}$ and the sample means μ_i^{data}; the ratio between the estimated model

rates $r_i^{iMEM_i}$ and the sample rates r_i^{data}; the ratio between the estimated second moments $S_{ij}^{iMEM_i}$ and the sample ones S_{ij}^{data}; the ratio between the estimated correlation coefficients $C_{ij}^{iMEM_i}$ and the sample ones C_{ij}^{data}; the scatter of pattern probabilities from the model (y-axis) vs pattern probabilities estimated from the data (x-axis), in a log-log scale (the solid line represents equality). Inside each scatter plots we also report the Kullback-Leibler divergence of the data from the model, and for iMEM$_2$ the information ratio $\mathcal{R}_2$.

A.4.2 Fit to independent data

As shown in Figure A.1, fitting iMEM$_1$to independent data leads to underestimate the means μ_i and therefore to overestimate the rates r_i (which have opposite sign, see Eq. ...). The second moments S_{ij} are also slightly underestimated for $i \neq j$. [2] However, the pattern probabilities are well estimated for patterns of size up to 8 (right-most panel). Due to the the finite size of the sample, pattern probabilities estimated from data have a resolution of $1/n = 10^{-6}$, which explains the deviation from the solid line for these pattern sizes. The high performance is captured by the Kullback-Leibler divergence $D_{KL} = 0.006$.

Model iMEM$_2$ is a generalization of iMEM$_1$. If the latter is the true model (as in this case), the first is expected to collapse to iMEM$_1$, i.e. $J_{ij} = 0\ \forall i \neq j$. However, we see from Figure A.1, second line, fourth plot, that the employed GIS procedure is not able to correctly determine the coefficients and assigns positive values to many of them, overfitting the positive correlations observed in the data due to noise. This, in conjunction with overestimated rates (second plot from the left) yields overestimated pattern probabilities for patterns of size 3 or higher, resulting in a worse performance than iMEM$_1$ ($D_{KL} = 0.016$). As a consequence, although no HOCs are present in the data, we have$\mathcal{R}_2 = 0.828$, significantly lower than 1. This warns against the use of this parameter as an indicator of information explained by pairwise vs. higher-order correlations (see Section A.5).

A.4.3 Fit to data with pairwise and higher-order correlations.

Figure A.2 shows the results obtained by fitting the two models to purely pairwise-correlated data. As expected, iMEM$_1$ provides a poor fit ($D_{KL} = 0.033$) because it does not capture the correlations, which it "converts" instead into over-estimated rates. Conversely, iMEM$_2$ provides an excellent fit ($D_{KL} = 0.001$). The index $\mathcal{R}_2$ is large ($> .99$)

Finally, Figure A.3 shows the results obtained by fitting the two models to data containing both pairwise and higher-order correlations. The pairwise correlations are ubiquitous ($C_{ij} = 0.01\ \forall i \neq j$), the HOCs are represented by a single group of $M = 5$ neurons spiking simultaneously $n_2 = 10^3$ times out of $n = 10^6$ total observations. This pattern is large compared to the dimension of the full system, reason for which it should be classified as highly significant by an HOC analysis. The independent model iMEM$_1$ yields a poor fit, as expected. At a first glance, iMEM$_2$ performs instead generally well ($D_{KL} = 0.003$). However, it heavily underestimates the probability of a few patterns: the injected SIP patterns (corresponding to the isolated orange dot in the bottom-right panel of Figure A.3) as well as of some of the patterns including it (the off-diagonal purple dots, corresponding to chance super-patterns of the injected SIP). Nonetheless, the information coefficient $\mathcal{R}_2$ is larger than 0.99, suggesting that the information about the full joint distribution of the data is carried almost completely by rates and pairwise correlations. This shows

[2] As for the correlation coefficients, note that the true ones used to generate the data are 0 for all $i \neq j$. The same holds, by definition, for the correlation coefficients $C_{ij}^{iMEM_i}$ iniMEM$_1$. The correlation coefficients computed over the data fluctuate around 10^{-4} due to stochasticity, resulting in the dark-blue entries in the small top-right panel in Figure A.1. Thus, the ratio $C_{ij}^{iMEM_i}/C_{ij}^{data}$ is 0 and not 1, but in this particular case this does not correspond to a bad performance.

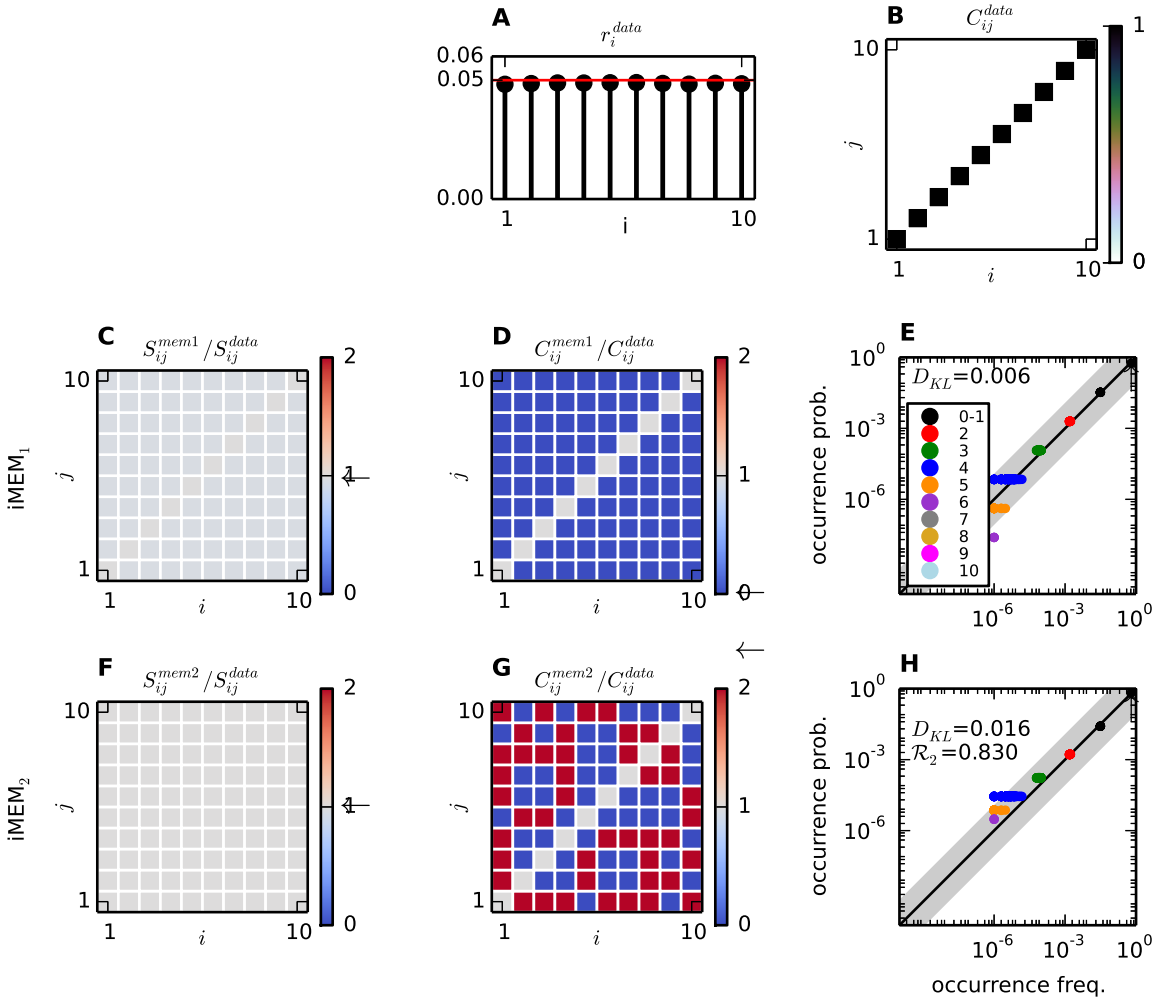

Figure A.1: **Fit to independent data.** Goodness of fit of iMEM$_1$ and iMEM$_2$ to data from $N = 10$ independent spike trains with stationary rates $r_i = 0.05$. **(A)** Firing rates evaluated from the generated data. The red line indicates the prescribed theoretical rates. **(B)** Matrix of pairwise correlation coefficients of the data. **(C-F)** Ratio of the second moments S_{ij} determined by models iMEM$_1$ (C) and iMEM$_2$ (F) over those observed in the data. The arrows in the colorbar indicate the average ratio over all non-diagonal elements of the matrix. **(D-G)** Ratio of the pairwise correlation coefficients C_{ij} determined by models iMEM$_1$ (D) and iMEM$_2$ (G) over those observed in the data. **(E-H)** Log-log scatter plot of the occurrence probability of network's states predicted by models iMEM$_1$ (E) and iMEM$_2$ (H) against the occurrence frequencies in the data. Different colors indicate states with a different number of active neurons (1s), from 0 to 10. The black straight line marks identical probabilities. The gray band around it indicates a mismatch of the two up to a factor of 10. The panels also report the values of the Kullback-Leibler divergence D_{KL}between the model's distribution and the true distribution. Panel H additionally shows the coefficient $\mathcal{R}_2 = 0.83$ of the model iMEM$_2$ fitted to the data.

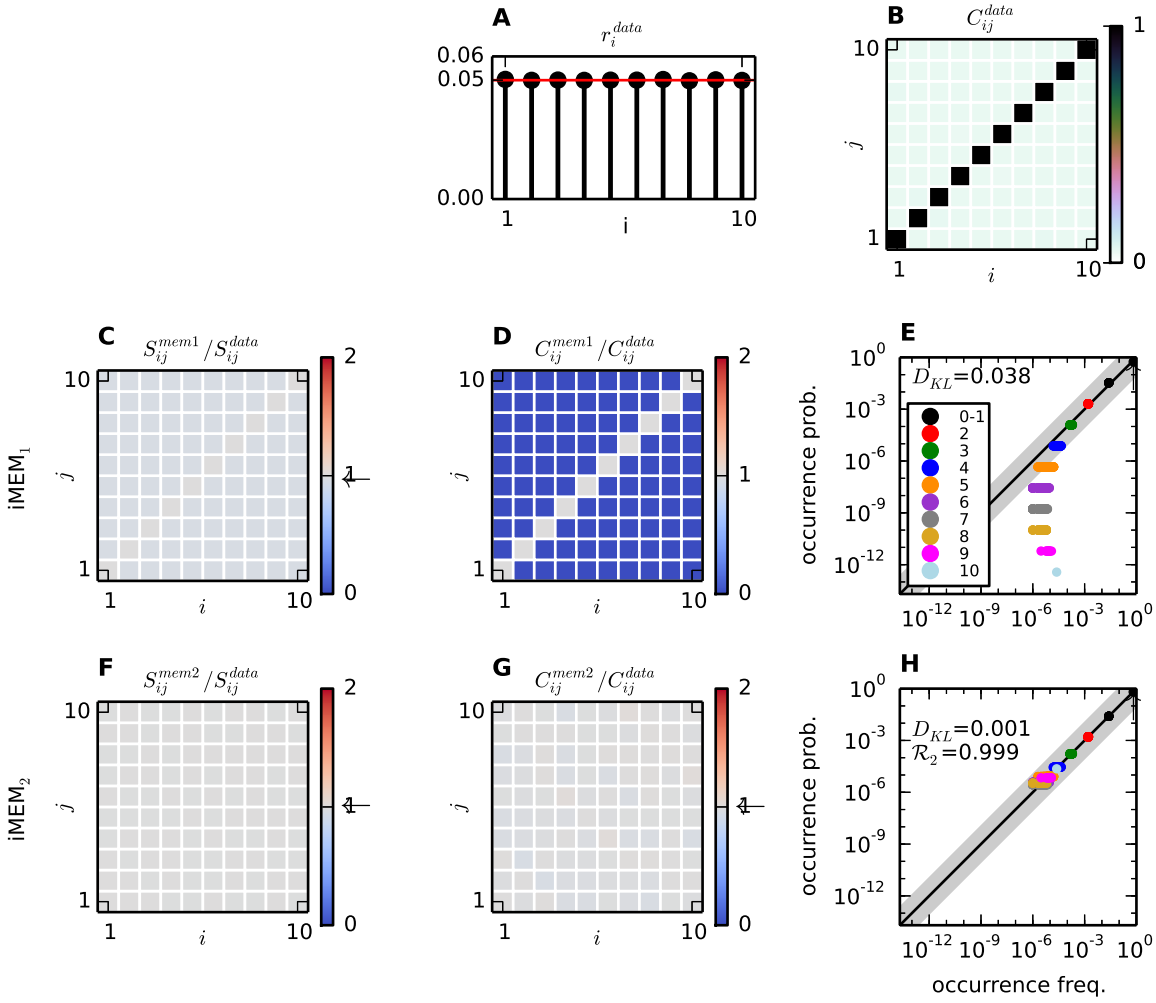

Figure A.2: **Fit to pairwise-correlated data.** Goodness of fit of iMEM$_1$ and iMEM$_2$ to data from $N = 10$ pairwise-correlated spike trains with stationary rates $r_i = 0.05$ and correlation coefficients $C_{ij} = 0.05$. Same as in Figure A.1.

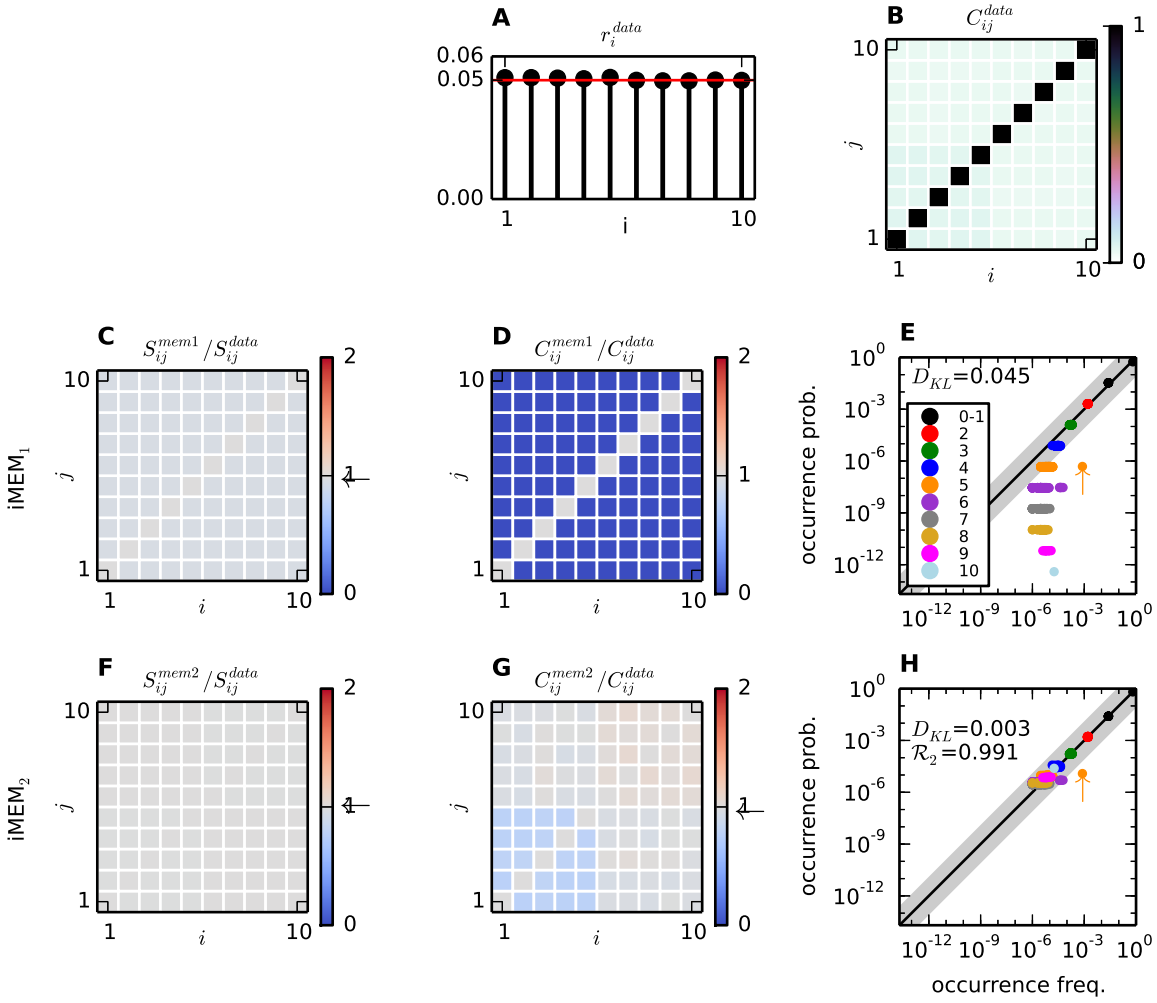

Figure A.3: **Fit to data with pairwise and higher-order correlations.** Goodness of fit of iMEM$_1$ and iMEM$_2$ to data from $N = 10$ pairwise-correlated spike trains with stationary rates $r_i = 0.05$ and correlation coefficients $C_{ij} = 0.05$. The first $M = 5$ neurons also exhibit synchronous activity in 10^3 of the total 10^6 bins, acting like a SIP. The panels are the same as in Figure A.1. The arrows in panels E, H indicate the dot corresponding to the inserted SIP pattern.

how misleading such a measure can be: although the "fraction of information" carried by HOCs is very small, one should not advocate for the absence of such correlations, which instead could exist and be indeed highly significant. The reason for this is that $\mathcal{R}_2$ is a ratio of entropy differences. Entropies are in turn sums of pattern probabilities multiplied by the logarithms thereof (see Equation (A.1)). If most but not all probabilities are well represented by the model, then $I^{(2)} \simeq I^{(N)}$ and $\mathcal{R}_2 \simeq 1$ despite the presence of a few highly mis-estimated probabilities. Increasing the SIP size to 8 and 10 (and therefore the order of correlation, up to a phenomenon involving all neurons) yielded information coefficients $\mathcal{R}_2 = 0.998$ and $\mathcal{R}_2 \simeq 1$ respectively (not shown), demonstrating the insensitivity of this measure to HOCs.

A.5 Discussion

Schneidman et al. (2003) proposed to quantify the relevance of pairwise vs higher-order correlations by the measure $\mathcal{R}_2$ defined in Equation (A.8). A value of $\mathcal{R}_2$ close to 1 would correspond to pairwise correlations being able to capture the full correlation structure in the data. A value lower than 1 would correspond to the presence of relevant HOCs. The measure was employed in Schneidman et al. (2006), Shlens et al. (2006) and Tang et al. (2008), who reported in all cases values close to 1 and thus allegedly negligible HOCs.

In Section A.4 we investigated the behavior of the measure $\mathcal{R}_2$ in the presence of independent, pairwise-correlated and HOC data. By doing so we were able to show that this index quantifies the relative influence of pairwise correlations on the probability distribution of spike patterns, but does not allow to draw conclusion on the statistical significance of possible HOCs. Indeed, we found scenarios where $\mathcal{R}_2$ is very close to 1 ($\mathcal{R}_2 \geq 0.99$) despite the presence of higher-order patterns of synchronous spikes with very low p-values. We interpret this dichotomy as follows: due to the presence of hundreds of possible patterns, the existence of individual significant patterns does not sensibly affect the overall pattern probability distribution, and likewise the value of $\mathcal{R}_2$. $\mathcal{R}_2$ can be used to quantify the influence of neuronal assemblies on the full probability distribution, but not to conclude on their presence.

Furthermore, iMEM$_2$ tends to introduce pairwise correlations when fitted to independent data (see Figure A.1). The parameters found by fitting the model via GIS (see Section A.2.5) reproduce the first and second moments reasonably well. However, the error in rates and correlation coefficients, related to the moments via Equations (A.9) and (A.10), is amplified. Similarly, pairwise correlations are properly estimated when strong enough ($C_{ij} = 0.05\ \forall i \neq j$; Figure A.2) but overestimated when weaker than that (true correlation coefficients $C_{ij} = 0.01$ are estimated to be ≈ 0.044). Raw moments are well estimated instead. Speculating that this may be a problem resulting from the fitting procedure, we modified it in two alternative ways: first, by replacing the factor $\frac{1}{k+1}$ in Equation (A.12) by $\frac{1}{\sqrt{1+k}}$, thus slowing down the vanishing corrections. Second, by replacing in the same equation the difference $\hat{f}_i^{(k)} - \hat{f}_i$ between obtained and desired moments at each step k with the difference between the obtained and desired correlation coefficients. None of the two approaches worked.

These results show that iMEM$_2$ is over-sensitive to low or missing correlations, and suggests some care before blindly fitting it to data characterized by weak sample correlation coefficients. In general, one should check that the sample correlations (and not just the moments) are well reproduced.

Despite these challenges and short-comings, maximum entropy modeling provides a powerful and interesting approach to fit parallel spike train data to a stochastic model. In

this appendix we highlighted the inefficacy of the measure $\mathcal{R}_2$ in representing the relevance of HOCs in MPST data. In Chapter 1 we also discussed that the absence of statistically significance binary vector states does not necessarily imply the absence of statistically significant patterns of synchronous spikes. Indeed, the latter may be present, however diluted in a number of different but overlapping states characterized by additional chance spikes from other neurons. Nevertheless, maximum entropy models may still be useful in HOC analyses. For instance, one may fit the data to iMEM_2 and then use the model to generate surrogate data with the same average firing rates and pairwise correlations. Then, one may use SPADE/ASSET (see Chapter 2 and Chapter 4, respectively) to investigate the data for HOCs/SSEs which are not explained by pairwise correlations.

Basic Notation

Abbreviations:

ASSET	analysis of sequences of synchronous events
CCH	cross-correlation histogram
CFIS	closed frequent item set
cdf	cumulative distribution function
COCONAD	continuous-time closed neuronal assembly detection
CPP	compound Poisson process
DS	diagonal structure
FIM	frequent item set mining
FIS	frequent item set
FN	false negative
FP	false positive
HOC	higher-order correlation
ISI	inter-spike interval
MPST	massively parallel spike trains
pdf	probability density function
pmf	probability mass function
PSF	pattern spectrum filtering
PSR	pattern set reduction
PSTH	peri-stimulus time historam
SIP	single interaction process
SNR	signal-to-noise ratio
SSE(s)	sequence(s) of synchronous events
STP	spatio-temporal pattern
SPADE	spike pattern detection and evaluation
UE	unitary event

Frequent mathematical symbols:

$:=$	defined as or denoted by
$\equiv$	identically equal to
$\approx$	approximately equal to
$\sim$	distributed as
H_0	null hypothesis
α	significance threshold of a statistical test
N	total number of neurons (spike trains) in a data set
c	pattern support (number of pattern occurrences)
z	pattern size (number of composing neurons)
$\langle z, c \rangle$	pattern signature

Physical quantities:

(m)s	(milli)second
(k)Hz	(kilo)hertz
μm/ mm	micro- / milli-meter
N	Newton
(m)V	(milli)volt

Acknowledgments

This work is the result of the collaboration with an incredible variety of people I had the luck to meet in the past four years. And of the bonds that lasted with those I knew from before. Here I wish to express my gratitude towards you all, in almost brainstorm order.

First, I wish to thank Prof. Sonja Grün, who let me join her lab and introduced me to the field of computational neuroscience. You put your trust in me, letting me work on and even lead exciting projects, while always offering your help and sharing your thoughts and opinion. By always asking questions, you were reminding me and our crew that no one should feel ashamed to ask. And that the socratic awareness of one's own ignorance is the highest form of knowledge.

Alongside with Sonja, I wish to express my gratitude to my other supervisor, Michael. You guided me in science with incredible discreteness. If I had to choose a single quality of yours which I benefit the most from, I would pick your calmness in front of sudden troubles and chaos. You have been a friend before a teacher, a role that few can take. You managed this, despite the headache coming from all those "grün" vs "dunkel". Learning German with you was a lot of fun.

Speaking of which, I owe my German to a much larger number of people who taught me Hochdeutsch and Umgangsprache, and warned me against the temptations of Svebisch. The first was my mother Silvana, who introduced me to this language with amazing teaching skill despite more than twenty years had passed since her last job as a translator. And then all my colleagues and friends. I wish to mention Abigail, Sacha, Hannah, Lyuba, Claudia, Tom, Moritz, Bernd, Jochen and Janine for correcting my mistakes, and to Jakob, Jannis, Philipp and Max for real life lessons while playing kicker, which I better not repeat here. If abroad they repeatedly mistook me for a German as I was talking in English, it's also your fault. Abigail, Hannah and Paul, you also have the merit of having always corrected my English, regardless of whether the situation made this proper or not. I will always owe to your mercylessness.

To all these people, and to the INM-6/IAS-6 in general, I am indebted for countless scientific discussions in our beautifully organized regular Journal CLub and Lab Seminar and Group Meeting and Jülich-Marseille Colloquium and Personal Meeting and PhD meeting and Internal Retreat and INM-retreat. And discussions in the corridors and in the kitchen. Your enthusiasm was contagious. Among you all, I wish to remember Paul: you have been the personification of a contagious enthusiasm. Also, thanks to Luca and Serge for broadening, with your different mindset, my interests and perspective, and for piling up interesting papers on my desk. I am as grateful to my office mates Lyuba and Andrei for sharing ideas and opinions on a daily basis and for trying to instill some notions of biology in this mathematical mind of mine.

Moritz, Luca and Vahid, I dearly value your appreciation and interest for some work of mine I was directing to the trash bin. I owe you the appendix of this thesis and a couple of ideas for future work. I intended to harvest these ideas with you much further, but time runs fast. Hopefully we will keep collaborating in the future.

Science is impossible without a properly organized working environment. Janine, Martina, Stefi, Joel and lastly Petra made my science possible and smooth, assisting me and all the lab way beyond their responsibilities. Bernd, Jochen, Alper and Tiziano provided incredible assistance and were always there by any computer crush, cluster problem, IT disaster in our lab, be it during working time, weekend, holiday or late at night. I wish to tell you that your daily administrative, technical and ultimately scientific work had a higher impact on my science than all equations I have written.

Science also comes out of thoughts. And thoughts come out of good food. Thanks to Jochen, Claudia, Tom, Espen, Hannah, Abigail, Julia, Luca and Robin, and other occasional lunch buddies, for sharing tasty food together. If it's true that you are what you eat, you made me a better person.

Thanks to Hannah, Philipp and Lyuba for letting me share with them so many thoughts and personal experiences, giving me good tips, transforming dramas into fun and reacting to obscure, unspekable secrets with a bored "That's it? So what?".

Time to move out of the lab. Being a (certainly privileged) immigrant in Germany was challenging, but amusing. I felt welcomed since the very first day. I am grateful to all who contributed to create this atmosphere. Thanks to all the friends I met along the way, starting from the very first one, Dima, who offered me shelter when I had just arrived. Vinh, you have been a great flatmate. Four years under the same roof officially make you the person who could cope with me for the longest time here in Germany. Thanks for the meals we shared occasionally, for all the chatting and exchange of opinions, and for saving our place from burning down a couple of times.

Occasionally, you meet someone you get to go out with only for a short time. Yet, the bond grows so fast you have the feeling it will last. Among others, thanks to Henri, Raffaele, Ariadna, Mika, Philipp, Hannah, Serge, Andrei, Sacha, Milena, Pino and Martin for countless evenings spent chatting about serious and stupid things, and for hosting me at your place, be it in Cologne, Rome, London, Okayama or Bremen. For sharing your Albanian, Roman, Mexican, Japanese, German, Russian, Dutch and Saarlandische culture with mine. For being, each in your own way, just so amazing. A special, tender "thank you" to kleined Martin for coping with me a few times more than necessary and for being willing to do so next.

I am, and will always be, deeply grateful to those who coped with my departure from Italy and stayed. A special mention to Elisa, Luisa, Chiara, Alessandro, Federica and Riccardo, who all kept contact even when I didn't reply, overwhelmed by my new life. Four years were a hard test for all of us. Thanks for managing it together, and sorry for trying your patience so hard. I dearly value the enduring friendship with Marina and Paolo too: our discussions around your dinner tab and amidst the sands of Sahara helped me not to lose connection with the world out there. Elsa and Vittorio, you always kept asking me about my life and have always waited for me to come back with tasty polenta and bagna cauda: your dearness is an ingredient that made your dishes unique.

I'd also like to put in words the love, the help, the care, the constant presence of my family. This thesis, the papers I wrote, my achievements would have not been possible without your moral but also the very practical support. My grandmother Nonna Ada was always so interested in knowing whether I finally found those "worms" in the brain: coming back and hearing you asking me what I was *really* doing was a priceless reward after months of hard work.

Finally, a few words to my family. Thanks for everything mom. Thanks for teaching me German long before I even thought of moving to Germany, which made my life here so easier and happier. And for the countless other things you did for me always. One understands always too late. Thank you brother for sharing experiences and for making

me doubt and think about the future. Thanks to you, the choice I took now seems much more mature than it was before. Having a twin, and you being the one, is the most awesome thing that could have happened to me. And thank you dad for providing me with immense support during hard times of deadlines which seemed impossible to meet. Analyzing situations with you is always so worth it. Thank you for your wisdom and for being such a good friend. All of you were present every single day, no matter the distance. Just be aware that I wouldn't be here, at this stage of my life, if it weren't for the three of you. Just know that I am happy.

To all of you: I don't know yet where my next steps will bring me. But I hope we will find out together.

References

Abeles, M. (1982a). *Local Cortical Circuits: An Electrophysiological Study*. Studies of Brain Function. Springer-Verlag, Berlin, Heidelberg, New York.

Abeles, M. (1982b). Role of cortical neuron: integrator or coincidence detector? *Israel Journal of Medical Sciences*, 18:83–92.

Abeles, M. (1991). *Corticonics: Neural Circuits of the Cerebral Cortex*. Cambridge University Press, Cambridge, 1st edition.

Abeles, M., Bergman, H., Margalit, E., and Vaadia, E. (1993). Spatiotemporal firing patterns in the frontal cortex of behaving monkeys. *Journal of Neurophysiology*, 70(4):1629–1638.

Abeles, M. and Gerstein, G. L. (1988). Detecting Spatiotemporal Firing Patterns Among Simultaneously Recorded Single Neurons. *Journal of Neurophysiology*, 60(3):909–924.

Adrian, E. (1926a). The impulses produced by sensory nerve-endings: Part 4. impulses from pain receptors. *Journal of Physiology*, 62(1):33–51.

Adrian, E. (1926b). The impulses produced by sensory nerve endings: Part i. *Journal of Physiology*, 61(1):49–72.

Adrian, E. and Zotterman, Y. (1926). The impulses produced by sensory nerve endings: Part 3. impulses set up by touch and pressure. *Journal of Physiology*, 61(4):465–83.

Aertsen, A., Gerstein, G., Habib, M., and Palm, G. (1989). Dynamics of neuronal firing correlation: modulation of "effective connectivity". *Journal of Neurophysiology*, 61(5):900–17.

Ahissar, M., Sosnik, R., and Haidarliu, S. (2000). Transformation from temporal to rate coding in somatosensory thalamocortical pathway. *Nature*, 406:302–306.

Ainsworth, M., Lee, S., Cunningham, M. O., Traub, R. D., Kopell, N. J., and Whittington, M. A. (2012). Rates and rhythms: A synergistic view of frequency and temporal coding in neuronal networks. *Neuron*, 75(4):572–83.

Arieli, A., Shoham, D., Hildesheim, R., and Grinvald, A. (1995). Coherent spatiotemporal patterns of ongoing activity revealed by real-time optical imaging coupled with single-unit recording in the cat visual cortex. *Journal of Neurophysiology*, 73(5):2072–2093.

Baker, S. N. and Lemon, R. (2000). Precise spatiotemporal repeating patterns in monkey primary and supplementary motor areas occur at chance level. *Journal of Neurophysiology*, 84:1770–1780.

Baker, S. N., Spinks, R., Jackson, A., and Lemon, R. N. (2001). Synchronization in monkey motor cortex during a precision grip task. I. Task-dependent modulation in single-unit synchrony. *Journal of Neurophysiology*, 85(2):869–85.

Barreiro, A. K., Gjorgjieva, J., Rieke, F., and Shea-Brown, E. (2014). When do microcircuits produce beyond-pairwise correlations? *Frontiers in Computational Neuroscience*, 8(10).

Beggs, J. M. and Plenz, D. (2004). Neuronal avalanches are diverse and precise activity patterns that are stable for many hours in cortical slice cultures. *The Journal of Neuroscience*, 24(22):5216–5229.

Bender, V. A., Bender, K. J., Brasier, D. J., and Feldman, D. E. (2006). Two coincidence detectors for spike timing-dependent plasticity in somatosensory cortex. *The Journal of Neuroscience*, 26(16):4166–77.

Berger, D., Borgelt, C., Louis, S., Morrison, A., and Grün, S. (2010). Efficient identification of assembly neurons within massively parallel spike trains. *Comput. Intell. and Neurosci.*, 2010:Article ID 439648.

Berger, D., Warren, D., Normann, R., Arieli, A., and Grün, S. (2007). Spatially organized spike correlation in cat visual cortex. *Neurocomputing*, 70(10-12):2112–2116.

Bienenstock, E. (1995). A model of neocortex. *Network*, 6:179–224.

Bliss, T. V. P. and Lomo, T. (1973). Long-lasting potentation of synaptic transmission in the dendate area of anaesthetized rabbit following stimulation of the perforant path. *Journal of Physiology*, 232:331–356.

Borgelt, C. (2012). Frequent item set mining. In *Wiley Interdisciplinary Reviews (WIREs): Data Mining and Knowledge Discovery*, volume 2, pages 437–456. J. Wiley & Sons, Chichester, United Kingdom. doi:10.1002/widm.1074.

Borgelt, C. and Picado-Muiño, D. (2014). Simple pattern spectrum estimation for fast pattern filtering with coconad. In Blockeel, H., van Leeuwen, M., and Vinciotti, V., editors, *Advances in Intelligent Data Analysis XIII*, volume 8819, pages 37–48. Springer International Publishing.

Braitenberg, V. and Schüz, A. (1991). *Anatomy of the Cortex: Statistics and Geometry*. Springer-Verlag, Berlin, Heidelberg, New York.

Braitenberg, V. and Schüz, A. (1998). *Cortex: Statistics and Geometry of Neuronal Connectivity*. Springer-Verlag, Berlin, 2nd edition.

Broderick, T., Dudik, M., Tkacik, G., Schapire, R., and Bialek, W. (2007). Faster solutions of the inverse pairwise ising problem. *arXiv*, page 0712.2437.

Buzsaki, G. (2004). Large-scale recording of neuronal ensembles. *Nature Neuroscience*, 7(5):446–451.

Carr, C. E. (2004). Timing is everything: organization of timing circuits in auditory and electrical sensory systems. *Journal of Comparative Neurology*, 472:131–133.

Confais, J., Kilavik, B., Ponce-Alvarez, A., and A, R. (2012). On the anticipatory pre-cue activity in motor cortex. *The Journal of Neuroscience*, 32:15359–68.

Darroch, J. N. and Ratcliff, D. (1972). Generalized iterative scaling for log-linear models. *The Annals of Mathematical Statistics*, 43(5):1470–1480.

Dayhoff, J. and Gerstein, G. (1983a). Favored patterns in spike trains. i. detection. *Journal of Neurophysiology*, 49(6):1334–48.

Dayhoff, J. E. and Gerstein, G. L. (1983b). Favored patterns in spike trains. II. application. *Journal of Neurophysiology*, 49(6):1349–1363.

De Gruijl, J., Hoogland, T., and De Zeeuw, C. (2014). Behavioral correlates of complex spike synchrony in cerebellar microzones. *The Journal of Neuroscience*, 34(27):8937–47.

De la Rocha, J., Doiron, B., Shea-Brown, E., Kresimir, J., and Reyes, A. (2007). Correlation between neural spike trains increases with firing rate. *Nature*, 448(16):802–807.

Denker, M., Riehle, A., Diesmann, M., and Grün, S. (2010). Estimating the contribution of assembly activity to cortical dynamics from spike and population measures. *Journal of Computational Neuroscience*, 29(3):599–613.

Diba, K., Amarasingham, A., Mizuseki, K., and Buzsàki, G. (2014). Millisecond timescale synchrony among hippocampal neurons. *The Journal of Neuroscience*, 34(45):14984–14994.

Diesmann, M., Gewaltig, M.-O., and Aertsen, A. (1996). Characterization of synfire activity by propagating 'pulse packets'. In Bower, J. M., editor, *Computational Neuroscience: Trends in Research*, pages 59–64, San Diego. Academic Press.

Diesmann, M., Gewaltig, M.-O., and Aertsen, A. (1999). Stable propagation of synchronous spiking in cortical neural networks. *Nature*, 402(6761):529–533.

Edelman, G. M. (1993). Neural darwinism: selection and reentrant signaling in higher brain function. *Neuron*, 10(2):115–25.

Eggermont, J. (2015). Animal models of auditory temporal processing. *International Journal of Psychophysiology*, 95(2):202–215.

Engel, A. K., König, P., Schillen, T. B., and Singer, W. (1992). Temporal coding in the visual cortex: new vistas on integration in the nervous system. *Trends in Neurosciences*, 15(6):218–226.

Ester, M., Kriegel, H.-P., Sander, J., and Xu, X. (1996). A density-based algorithm for discovering clusters in large spatial databases with noise. *Proceedings of 2nd International Conference on Knowledge Discovery and Data Mining (KDD-96)*, pages 226–231.

Farkhooi, F., Muller, E., and Nawrot, M. P. (2011). Adaptation reduces variability of the neuronal population code. *Physical Review E*, 83(5):050905.

Farkhooi, F., Strube-Bloss, M. F., and Nawrot, M. P. (2009). Serial correlation in neural spike trains: Experimental evidence, stochastic modeling, and single neuron variability. *Physical Review E*, 79(2):021905.

Feldman, D. E. (2000). Timing-based LTP and LTD at vertical inputs to layer II/III pyramidal cells in rat barrel cortex. *Neuron*, 27:45–56.

Fino, E., Paille, V., Cui, Y., Morera-Herreras, T., Deniau, J., and Venance, L. (2010). Distinct coincidence detectors govern the corticostriatal spike timing-dependent plasticity. *Journal of Physiology*, 588(16):3045–3062.

Fujisawa, S., Amarasingham, A., Harrison, M. T., and Buzsáki, G. (2008). Behavior-dependent short-term assembly dynamics in the medial prefrontal cortex. *Nature Neuroscience*, 11(7):823–833.

Gansel, K. S. and Singer, W. (2012). Detecting multineuronal temporal patterns in parallel spike trains. *Frontiers in Neuroinformatics*, 6.

Gautrais, J. and Thorpe, S. (1998). Rate coding versus temporal order coding: a theoretical approach. *Biosystems*, 48:57–65.

Georgopoulos, A., Schwartz, A., and Kettner, R. (1986). Neuronal population coding of movement direction. *Science*, 4771(233):1416–1419.

Gerstein, G. L., Williams, E. R., Diesmann, M., Grün, S., and Trengove, C. (2012). Detecting synfire chains in parallel spike data. *Journal of Neuroscience Methods*, 206:54 – 64.

Gerstner, W., Kreiter, A. K., Markram, H., and Herz, V. M. (1997). Neural codes: Firing rates and beyond. *Proceedings of the National Academy of Sciences*, 94:12740–12741.

GL Gerstein, DH Perkel, J. D. (1985). Cooperative firing activity in simultaneously recorded populations of neurons: detection and measurement. *The Journal of Neuroscience*, 5(4):881–889.

Goethals, M. D. (2010). Data mining and knowledge discovery handbook. chapter Frequent Set Mining, pages 321–338. Springer Verlag, Berlin, 2nd edition. doi: 10.1007/978-0-387-09823-4_16.

Gray, C. M., König, P., Engel, A. K., and Singer, W. (1989). Oscillatory responses in cat visual cortex exhibit inter-columnar synchronization which reflects global stimulus properties. *Nature*, 338:334–337.

Grün, S. (2009). Data-driven significance estimation of precise spike correlation. *Journal of Neurophysiology*, 101(3):1126–1140.

Grün, S., Abeles, M., and Diesmann, M. (2008). Impact of higher-order correlations on coincidence distributions of massively parallel data. In *Lecture Notes in Computer Science, 'Dynamic Brain - from Neural Spikes to Behaviors'*, volume 5286, pages 96–114.

Grün, S., Diesmann, M., and Aertsen, A. (2002a). 'Unitary Events' in multiple single-neuron spiking activity. I. Detection and significance. *Neural Computation*, 14(1):43–80.

Grün, S., Diesmann, M., and Aertsen, A. (2002b). 'Unitary Events' in multiple single-neuron spiking activity. II. Non-Stationary data. *Neural Computation*, 14(1):81–119.

Grün, S., Diesmann, M., and Aertsen, A. (2002c). 'Unitary Events' in multiple single-neuron spiking activity. II. Non-Stationary data. *Neural Computation*, 14(1):81–119.

Grün, S., Diesmann, M., Grammont, F., Riehle, A., and Aertsen, A. (1999a). Detecting unitary events without discretization of time. *Journal of Neuroscience Methods*, 94(1):67–79.

Grün, S., Diesmann, M., Grammont, F., Riehle, A., and Aertsen, A. (1999b). Detecting unitary events without discretization of time. *Journal of Neuroscience Methods*, 94(1):67–79.

Grün, S., Riehle, A., and Diesmann, M. (2003). Effect of cross-trial nonstationarity on joint-spike events. *Biological Cybernetics*, 88(5):335–351.

Grün, S. and Rotter, S., editors (2010). *Analysis of Parallel Spike Trains*. Springer.

Gütig, R., Aertsen, A., and Rotter, S. (2003). Analysis of higher-order neuronal interactions based on conditional inference. *Biological Cybernetics*, 88(5):352–359.

Hao, Y., Riehle, A., and Brochier, T. (2015). Mapping the horizontal cortical connections in primate motor cortex using single pulse icms. In *Society for Neuroscience Abstracts*, page 244.11.

Harris, K. (2005). Neural signatures of cell assembly organization. *Nature Reviews Neuroscience*, 5(6):339–407.

Harvey, M. A., Saal, H. P., III, J. F. D., and Bensmaia, S. J. (2013). Multiplexing stimulus information through rate and temporal codes in primate somatosensory cortex. *PLoS Biology*, 11(5):e1001558.

Hatsopoulos, N., Joshi, J., and O'Leary, J. G. (2004). Decoding continuous and discrete motor behaviors using motor and premotor cortical ensembles. *Journal of Neurophysiology*, 92:1165–1174.

Hebb, D. O. (1949). *The organization of behavior: A neuropsychological theory*. John Wiley & Sons, New York.

Helias, M., Tetzlaff, T., and Diesmann, M. (2013). Echoes in correlated neural systems. *New Journal of Physics*, 15:023002.

Hodgkin, A. L. and Huxley, A. F. (1952a). Currents carried by sodium and potassium ions through the membrane of the giant axon of *Loligo*. *Journal of Physiology*, 116:449–472.

Hodgkin, A. L. and Huxley, A. F. (1952b). A quantitative description of membrane current and its application to conduction and excitation in nerve. *Journal of Physiology*, 117:500–544.

Hong, S., Ratté, S., Prescott, S. A., and De Schutter, E. (2012). Single neuron firing properties impact correlation-based population coding. *The Journal of Neuroscience*, 32(4):1413–1428.

Hosaka, R., Araki, O., and Ikeguchi, T. (2008). STDP provides the substrate for igniting synfire chains by spatiotemporal input patterns. *Neural Computation*, 20:415–435.

Hu, E. and Bloomfield, S. (2003). Gap junctional coupling underlies the short-latency spike synchrony of retinal alpha ganglion cells. *The Journal of Neuroscience*, 23(17):6768–77.

Hu, H. and Agmon, A. (2015). Properties of precise firing synchrony between synaptically coupled cortical interneurons depend on their mode of coupling. *Journal of Neurophysiology*, 114(1):624–37.

Ikegaya, Y., Aaron, G., Cossart, R., Aronov, D., Lampl, I., Ferster, D., and Yuste, R. (2004). Synfire chains and cortical songs: temporal modules of cortical activity. *Science*, 5670(304):559–564.

Ito, J., Maldonado, P., Singer, W., and Grün, S. (2011). Saccade-related modulations of neuronal excitability support synchrony of visually elicited spikes. *Cerebral Cortex*, 21(11):2482–2497.

Izhikevich, E. M. (2006). Polychronization: Computation with spikes. *Neural Computation*, 18:245–282.

Jackson, A., Gee, V., Baker, S., and Lemon, R. (2003). Synchrony between neurons with similar muscle fields in monkey motor cortex. *Neuron*, 1(38):115–125.

Jaynes, E. T. (1957). Information theory and statistical mechanics. *The Physical Review*, 106(4):620–630.

Kandel, E. R., Schwartz, J. H., and Jessell, T. M. (1991). *Principles of neural science*. McGraw-Hill, New York, 3 edition.

Kass, R. E., Kelly, R. C., and Loh, W.-L. (2011). Assessment of synchrony in multiple neural spike trains using loglinear point process models. *The Annals of Applied Statistics*, 5(2B):1262–1292.

Kaufman, L. and Rousseeuw, P. (1987). Clustering by means of medoids. In Dodge, Y., editor, *Statistical Data Analysis Based on the L1 Norm and Related Methods*, page 405:416. North-Holland.

Kelly, R. C. and Kass, R. E. (2012). A framework for evaluating pairwise and multiway synchrony among stimulus-driven neurons. *Journal of Neural Computation*, 24(8):2007–2032.

Kilavik, B. E., Roux, S., Ponce-Alvarez, A., Confais, J., Gruen, S., and Riehle, A. (2009). Long-term modifications in motor cortical dynamics induced by intensive practice. *The Journal of Neuroscience*, 29:12653–12663.

König, P., Engel, A. K., and Singer, W. (1996). Integrator or coincidence detector? The role of the cortical neuron revisited. *Trends in Neurosciences*, 19(4):130–137.

Kuhn, A., Aertsen, A., and Rotter, S. (2003a). Higher-order statistics of input ensembles and the response of simple model neurons. *Neural Computation*, 15(1):67–101.

Kuhn, A., Aertsen, A., and Rotter, S. (2003b). Higher-order statistics of input ensembles and the response of simple model neurons. *Neural Computation*, 15:67–101.

Kumar, A., Rotter, S., and Aertsen, A. (2010). Spiking activity propagation in neuronal networks: reconciling different perspectives on neural coding. *Nature Neuroscience*, 11:615–627.

Kwan, H. C., Murphy, J. T., and Wong, Y. C. (1987). Interaction between neurons in precentral cortical zones controlling different joints. *Brain Research*, 400:259–269.

Lamme, V. and Spekreijse, H. (1998). Neuronal synchrony does not represent texture segregation. *Nature*, 396(6709):362–366.

Lang, E. J. (2002). Gabaergic and glutamatergic modulation of spontaneous and motor-cortex complex spike activity. *Journal of Neurophysiology*, 87(4):1993–2008.

Le Cam, L. (1963). On the distribution of sums of independent random variables. In Neyman, J. and Cam, L. M. L., editors, *Bernoulli, Bayes, Laplace: Proceeding of an International Research Seminar*, pages 179–202. New York: Springer-Verlag.

Lømo, T. (2003). The discovery of long-term potentiation. *Philosophical Transaction of the Royal Society London, Serie B.*, 358:617–620.

London, M., Roth, A., Beeren, L., Häusser, M., and Latham, P. E. (2010). Sensitivity to perturbations in vivo implies high noise and suggests rate coding in cortex. *Nature Reviews Neuroscience*, 466(1):123–128.

Louis, S., Gerstein, G. L., Grün, S., and Diesmann, M. (2010). Surrogate spike train generation through dithering in operational time. *Frontiers in Computational Neuroscience*, 4(127).

Markram, H., Gerstner, W., and Sjöström, P. (2011). A history of spike-timing-dependent plasticity. *Frontiers in Synaptic Neuroscience*, 3(4).

Markram, H., Lübke, J., Frotscher, M., and Sakmann, B. (1997). Regulation of synaptic efficacy by coincidence of postsynaptic APs and EPSPs. *Science*, 275:213–215.

Martignon, L., Deco, G., Laskey, K., Diamond, M., Freiwald, W., and Vaadia, E. (2000). Neural coding: Higher-order temporal patterns in the neurostatistics of cell assemblies. *Neural Computation*, 12:2621–2653.

Martignon, L., von Hasseln, H., Grün, S., Aertsen, A., and Palm, G. (1995). Detecting higher-order interactions among the spiking events in a group of neurons. *Biological Cybernetics*, 73:69–81.

Martin, A. B. and von der Heydt, R. (2015). Spike synchrony reveals emergence of proto-objects in visual cortex. *The Journal of Neuroscience*, 37(15):6860–6870.

Massey, P. and Bashir, Z. (2007). Long-term depression: multiple forms and implications for brain function. *Trends in Neurosciences*, 30(4):176–184.

Milekovic, T., Truccolo, W., Grün, S., Riehle, A., and Brochier, T. (2015). Local field potentials in primate motor cortex encode grasp kinetic parameters. *Neuroimage*, 114:338–55.

Murthy, V. N. and Fetz, E. (1996). Synchronization of neurons during local field potential oscillations in sensorimotor cortex of awake monkeys. *Journal of Neurophysiology*, 76:3968–3982.

Nadasdy, Z., Hirase, H., Czurko, A., Csicsvari, J., and Buzsaki, G. (1999). Replay and time compression of recurring spike sequences in the hippocampus. *The Journal of Neuroscience*, 19(21):9497–9507.

Nakahara, H. and Amari, S. (2002). Information-geometric measure for neural spikes. *Neural Computation*, 10(14):2269–2316.

Nawrot, M., Aertsen, A., and Rotter, S. (1999). Single-trial estimation of neuronal firing rates: from single-neuron spike trains to population activity. *Journal of Neuroscience Methods*, 1(94):81–92.

Nawrot, M. P., Riehle, A., Aertsen, A., and Rotter, S. (2000). Spike count variability in motor cortical neurons. In *European Journal of Neuroscience*, volume 12, Supl. 11, page 506.

Ochiai, A. (1957). Zoogeographical studies on the soleoid fishes found japan and its neighboring regions. *Bulletin of the Japanese Society for the Science of Fish*, 22(9):526–530.

Ohiorhenuan, I. E., Mechler, F., Purpura, K. P., Schmid, A. M., Hu, Q., and Victor, J. D. (2010). Sparse coding and high-order correlations in fine-scale cortical networks. *Nature*, 466:617–621.

Pazienti, A., Maldonado, P. E., Diesmann, M., and Grün, S. (2008). Effectiveness of systematic spike dithering depends on the precision of cortical synchronization. *Brain Research*, 1225:39–46.

Perez-Orive, J., Bazhenov, M., and Laurent, G. (2004). Intrinsic and circuit properties favor coincidence detection for decoding oscillatory input. *The Journal of Neuroscience*, 24(26):6037–47.

Perkel, D. H., Gerstein, G. L., and Moore, G. P. (1967a). Neuronal spike trains and stochastic point processes. I. The single spike train. *Biophysical Journal*, 7(4):391–418.

Perkel, D. H., Gerstein, G. L., and Moore, G. P. (1967b). Neuronal spike trains and stochastic point processes. II. Simultaneous spike trains. *Biophysical Journal*, 7(4):419–440.

Picado-Muiño, D., Borgelt, C., Berger, D., Gerstein, G. L., and Grün, S. (2013). Finding neural assemblies with frequent item set mining. *Frontiers in Neuroinformatics*, 7(9). doi: 10.3389/fninf.2013.00009.

Pillow, J. W., Shlens, J., Paninski, L., Sher, A., Litke, A. M., Chichilnisky, E. J., and Simoncelli, E. P. (2008). Spatio-temporal correlations and visual signalling in a complete neuronal population. *Nature*, 454:995–999.

Pipa, G., Grün, S., and van Vreeswijk, C. (2013a). Impact of spike train autostructure on probability distribution of joint spike events. *Neural Computation*, 25(5):1123–63.

Pipa, G. and Munk, M. H. J. (2011). Higher order spike synchrony in prefrontal cortex during visual memory. *Frontiers in Computational Neuroscience*, 5(23).

Pipa, G., Riehle, A., and Grün, S. (2007). Validation of task-related excess of spike coincidences based on neuroxidence. *Neurocomputing*, 70:2064–2068. published online 2006; doi:10.1016/j.neucom.2006.10.142.

Pipa, G., van Vreeswijk, C., and Grün, S. (2013b). Impact of spike-train autostructure on probability distribution of joint spike events. *Neural Computation;*, 25(5):1123–63.

Pipa, G., Wheeler, D. W., Singer, W., and Nikolic, D. (2008). Neuroxidence: reliable and efficient analysis of an excess or deficiency of joint-spike events. *Journal of Neuroscience Methods*, 25:64 – 88.

Prescott, S. and Sejnowski, T. (2008). Spike-rate coding and spike-time coding are affected oppositely by different adaptation mechanisms. *The Journal of Neuroscience*, 28(50):13649–61.

Prut, Y., Vaadia, E., Bergman, H., Haalman, I., Hamutal, S., and Abeles, M. (1998). Spatiotemporal structure of cortical activity: Properties and behavioral relevance. *Journal of Neurophysiology*, 79(6):2857–2874.

Reed, J., Pouget, P., Qi, H., Zhou, Z., Bernard, M., Burish, M., Haitas, J., Bonds, A., and JH., K. (2008). Widespread spatial integration in primary somatosensory cortex. *Proceedings of the National Academy of Sciences of the United States of America*, 105(29):10233–7.

Reimer, I., Staude, B., Boucsein, C., and Rotter, S. (2013). A new method to infer higher-order spike correlations from membrane potentials. *Journal of Computational Neuroscience*, 35(2):169–86.

Reimer, I., Staude, B., Ehm, W., and Rotter, S. (2012). Modeling and analyzing higher-order correlations in non-poissonian spike trains. *Journal of Neuroscience Methods*, 208(1):18–33.

Riehle, A. (2005). Preparation for action: one of the key functions of the motor cortex. In Riehle, A. and Vaadia, E., editors, *Motor Cortex in Voluntary Movements: A Distributed System for Distributed Functions*, pages 213–240. CRC-Press, Boca Raton, FL.

Riehle, A., Grün, S., Diesmann, M., and Aertsen, A. (1997). Spike synchronization and rate modulation differentially involved in motor cortical function. *Science*, 278(5345):1950–1953.

Riehle, A. and Requin, J. (1989). Monkey primary motor and premotor cortex: single cell activity related to prior information about direction and extent of an intended movement. *Journal of Neurophysiology*, 61:534–549.

Riehle, A., Wirtssohn, S., Grün, S., and Brochier, T. (2013). Mapping the spatio-temporal structure of motor cortical lfp and spiking activities during reach-to-grasp movements. *Frontiers in Neural Circuits*, 7:48. doi: 10.3389/fncir.2013.00048.

Roelfsema, P. R., Lamme, V. A. F., and Spekreijse, H. (2004). Synchrony and covariation of firing rates in the primary visual cortex during contour grouping. *Nature Neuroscience*, 7(9):982–991.

Roudi, Y., Nirenberg, S., and Latham, P. E. (2009). Pairwise maximum entropy models for studying large biological systems: when they can work and when they canâĂŹt. *Plos Computational Biology*, 5(5):e1000380.

Roy, S. A. and Alloway, K. D. (2001). Coincidence detection or temporal integration? what the neurons in somatosensory cortex are doing. *The Journal of Neuroscience*, 21(7):2462–2473.

Rubino, D., Robbins, K. A., and Hatsopoulos, N. G. (2006). Propagating waves mediate information transfer in the motor cortex. *Nature Neuroscience*, 9:1549–1557.

Rudolph, M. and Destexhe, A. (2003). Tuning neocortical pyramidal neurons between integrators and coincidence detectors. *Journal of Computational Neuroscience*, 14:239–251.

Sakurai, Y. (1996). Hippocampal and neocortical cell assemblies encode memory processes for different types of stimuli in the rat. *The Journal of Neuroscience*, 16(8):2809–19.

Sakurai, Y. and Takahashi, S. (2006). Dynamic synchrony of firing in the monkey prefrontal cortex during working-memory tasks. *The Journal of Neuroscience*, 6(40):10141–10153.

Schneidman, E., Berry, M. J., Segev, R., and Bialek, W. (2006). Weak pairwise correlations imply strongly correlated network states in a neural population. *Nature*, 440:1007–1012.

Schneidman, E., Still, S., II, M. J. B., and Bialek, W. (2003). Network information and connected correlations. *Phisical Review Letters*, 91(23):238701.

Schrader, S., Grün, S., Diesmann, M., and Gerstein, G. (2008). Detecting synfire chain activity using massively parallel spike train recording. *Journal of Neurophysiology*, 100:2165–2176.

Schultze-Kraft, M., Diesmann, M., Gruen, S., and Helias, M. (2013). Noise suppression and surplus synchrony by coincidence detection. *PLoS Computational Biology*, 9(4):e1002904.

Schwarz, D. A., Lebedev, M. A., Hanson, T. L., Dimitrov, D. F., Lehew, G., Meloy, J., Rajangam, S., Subramanian, V., Ifft, P. J., Li, Z., Ramakrishnan, A., Tate, A., Zhuang, K. Z., and Nicolelis, M. A. L. (2014). Chronic, wireless recordings of large-scale brain activity in freely moving rhesus monkeys. *Nature Methods*, 11:670–676.

Seki, S. and Eggermont, J. J. (2003). Changes in spontaneous firing rate and neural synchrony in cat primary auditory cortex after localized tone-induced hearing loss. *Hearing Research*, 180(1-2):28–38.

Shadlen, M. N. and Movshon, A. J. (1999). Synchrony unbound: A critical evaluation of the temporal binding hypothesis. *Neuron*, 24:67–77.

Shadlen, M. N. and Newsome, W. T. (1994). Noise, neural codes and cortical organization. *Current Opinion in Neurobiology*, 4(4):569–579.

Shadlen, M. N. and Newsome, W. T. (1995). Is there a signal in the noise? *Current Opinion in Neurobiology*, 5:248–250.

Shannon, E. C. (1948). A mathematical theory of communication. *The Bell System Technical Journal*, 27:379–423.

Shea-Brown, E., Josic, K., de la Rocha, J., and Doiron, B. (2008). Correlation and synchrony transfer in integrate-and-fire neurons: basic properties and consequences for coding. *Physical Review Letters*, 100:108102.

Shimazaki, H., Amari, S.-i., Brown, E. N. B., and Grün, S. (2012). State-space analysis of time-varying higher-order spike correlation for multiple neural spike train data. *Plos Computational Biology*, 8(3):e1002385. doi:10.1371/journal.pcbi.1002385.

Shimazaki, H. and Shinomoto, S. (2007). A method for selecting the bin size of a time histogram. *Neural Computation*, 19(6):1503–1527.

Shimazaki, H. and Shinomoto, S. (2010). Kernel bandwidth optimization in spike rate estimation. *Journal of Computational Neuroscience*, 29:171–182.

Shinomoto, S., Sakai, Y., and Funahashi, S. (1999). The ornstein-uhlenbeck process does not reproduce spiking statistics of neurons in prefrontal cortex. *Neural Computation*, 11(4):935–951.

Shlens, J., Field, G. D., Gauthier, J. L., Grivich, M. I., Petrusca, D., Sher, A., Litke, A. M., and Chichilnisky, E. (2006). The structure of multi-neuron firing patterns in primate retina. *The Journal of Neuroscience*, 26(32):8254–8266.

Shmiel, T., Drori, R., Shmiel, O., Ben-Shaul, Y., Nadasdy, Z., Shemesh, M., Teicher, M., and Abeles, M. (2006). Temporally precise cortical firing patterns are associated with distinct action segments. *Journal of Neurophysiology*, 96(5):2645–2652.

Singer, W., Engel, A. K., Kreiter, A. K., Munk, M. H. J., Neuenschwander, S., and Roelfsema, P. R. (1997). Neuronal assemblies: necessity, signature and detectability. *Trends in Cognitive Sciences*, 1(7):252–261.

Sjöström, P., Turrigiano, G., and Nelson, S. (2001). Rate, timing, and cooperativity jointly determine cortical synaptic plasticity. *Neuron*, 32:1149–1164.

Smith, M. A. and Kohn, A. (2008). Spatial and temporal scales of neuronal correlation in primary visual cortex. *The Journal of Neuroscience*, 28(48):12591–12603.

Smith, W. and Fetz, E. (2009). Synaptic interactions between forelimb-related motor cortex neurons in behaving primates. *Journal of Neurophysiology*, 102:1026–1039.

Softky, W. R. (1995). Simple codes versus efficient codes. *Current Opininon in Neurobiology*, 5:239–247.

Staude, B., Grün, S., and Rotter, S. (2010a). Higher-order correlations in non-stationary parallel spike trains: statistical modeling and inference. *Frontiers in Computational Neuroscience*, 4:16.

Staude, B., Rotter, S., and Grün, S. (2008). Can spike coordination be differentiated from rate covariation? *Neural Computation*, 20:1973–1999.

Staude, B., Rotter, S., and Grün, S. (2010b). Cubic: cumulant based inference of higher-order correlations in massively parallel spike trains. *Journal of Computational Neuroscience*, 29:327 – 350. DOI 10.1007/s10827-009-0195-x.

Steinmetz, P., Roy, A., Fitzgerald, P., Hsiao, S., Johnson, K., and Niebur, E. (2000). Attention modulates synchronized neuronal firing in primate somatosensory cortex. *Nature*, 404(6774):187–190.

Stent, G. (1973). A physiological mechanism for Hebb's postulate of learning. *Proceedings of the National Academy of Sciences of the United States of America*, 70(4):997–1001.

Stuart, G. J. and Sakmann, B. (1994). Active propagation of somatic action potentials into neocortical pyramidal cell dendrites. *Nature*, 367(6458):69–72.

Takahashi, K., Kim, S., Coleman, T., Brown, K., Suminski, A., Best, M., and Hatsopoulos, N. (2015). Large-scale spatiotemporal spike patterning consistent with wave propagation in motor cortex. *Nature Communications*, 6:7169.

Tang, A., Jackson, D., Hobbs, J., Chen, W., Smith, J. L., Patel, H., Pietro, A., Petrusca, D., Grivich, M. I., Sher, A., Hottowy, P., Dabrowski, W., Litke, A. M., and Beggs, J. M. (2008). A maximum entropy model applied to spatial and temporal correlations from cortical networks in vitro. *The Journal of Neuroscience*, 28(2):505–518.

Tkacik, G., Schneidman, E., II, M. J. B., and Bialek, W. (2006). Ising models for networks of real neurons. *arXiv:q-bio.NC*.

Torre, E., Canova, C., Denker, M., Gerstein, G. L., Helias, M., and Grün, S. (2015a). ASSET: Analysis of sequences of synchronous events in massively parallel spike trains. submitted.

Torre, E., Picado-Muiño, D., Denker, M., Borgelt, C., and Grün, S. (2013). Statistical evaluation of synchronous spike patterns extracted by frequent item set mining. *Frontiers in Computational Neuroscience*, 7:132.

Torre, E., Quaglio, P., Denker, M., Brochier, T., Riehle, A., and Grün, S. (2015b). Emergence of synchronous spike patterns in monkey motor cortex during a delayed reach-to-grasp task. submitted.

Vaadia, E., Haalman, I., Abeles, M., Bergman, H., Prut, Y., Slovin, H., and Aertsen, A. (1995). Dynamics of neuronal interactions in monkey cortex in relation to behavioural events. *Nature*, 373(6514):515–518.

van der Togt, C., Kalitzin, S., Spekreijse, H., Lamme, V., and Supèr, H. (2006). Synchrony dynamics in monkey V1 predict success in visual detection. *Cerebral Cortex*, 16(1):136–148.

Van Rullen, R. and Thorpe, S. (2001). Rate coding versus temporal order coding: what the retinal ganglion cells tell the visual cortex. *Neural Computation*, 13(6):1255–83.

VanRullen, R., Guyonneau, R., and Thorpe, S. J. (2005). Spike times make sense. *Trends in Neurosciences*, 28(1):1–4.

Villa, A. E. and Abeles, M. (1990). Evidence for spatiotemporal firing patterns within the auditory thalamus of the cat. *Brain Research*, 509(2):325–327.

Villa, A. E. P., Tetko, I. V., Hyland, B., and Najem, A. (1999). Spatiotemporal activity patterns of rat cortical neurons predict responses in a conditioned task. *Proceedings of the National Academy of Sciences*, 96(3):1106–1111.

von der Malsburg, C. (1986). Am I thinking assemblies? In Palm, G. and Aertsen, A., editors, *Brain Theory*, pages 161–176. Springer-Verlag, Berlin.

Waddington, A., Appleby, P. A., de Kamps, M., and Cohen, N. (2011). Emergence of synfire chains with triphasic spike-time-dependent plasticity. *BMC Neuroscience*, 12(Suppl 1):P41.

Wise, S. P., Weinrich, M., and Mauritz, K. H. (1983). Motor aspects of cue-related neuronal activity in premotor cortex of the rhesus monkey. *Brain Research*, 260:301–305.

Zhang, L. I., Tao, H. W., Holt, C. E., Harris, W. A., and Poo, M. M. (1998). A critical window for cooperation and competition among developing retinotectal synapses. *Nature*, 395:37–44.