

**Statistical Evaluation of Dithering Methods
for Pattern Detection**

Peter Bouss

Master Thesis in Physics
presented to the

Faculty of Mathematics, Computer Science
and Natural Sciences

at the
RWTH Aachen University

Abstract

The study of neuronal activity patterns is one of the high-interest topics in neuroscientific research (Abeles (2010)). It is hypothesized that highly interconnected groups of neurons, so-called cell assemblies, can be building blocks for information processing (Hebb (1949); Abeles (1991); Harris (2005)). One of the effects of these cell assemblies' activation would be the formation of precisely timed arrangements of spikes emitted by the neurons participating in an assembly (Bienenstock (1995); Izhikevich (2006)). Experimental evidence for millisecond precise spiking activity is available for synchrony patterns (Riehle et al. (1997); Kilavik et al. (2009); Torre et al. (2016)) and for patterns with temporal delays (Prut et al. (1998); Villa and Abeles (1990); Russo and Durstewitz (2017)).

One of the algorithms developed in recent years to verify this hypothesis is called SPADE (Torre et al. (2013, 2016); Quaglio et al. (2017); Quaglio (2019); Stella et al. (2019)), which stands for Spike PAttern Detection and Evaluation. This method compares electrophysiological data against a null hypothesis. Due to the complexity of the spiking recordings, this null hypothesis is not determined analytically but obtained by a Monte Carlo approach, i.e., by opportunely generated surrogates (Torre et al. (2013)).

This thesis aims to compare different surrogate methods, specifically regarding their application to form a statistically robust null hypothesis. Thereby, we refer to a sub-category of surrogate approaches called dithering methods (Louis et al. (2010b)), whose particular feature is that each spike is displaced individually within a small window.

We proceed in such a way that first, we summarize the necessary neuroscientific background in chapter 1. After a brief explanation of the composition of neurons, we will introduce action potentials, or spikes, that serve to propagate information. We contrast the two standard hypotheses for neural coding and provide the reasons behind the expectation to find spatio-temporal patterns.

The second chapter gives an overview of the theory of statistical point processes used to model spike trains, i.e., sequences of spike times. We introduce the common statistical measures used to describe electrophysiological recordings. In particular, we focus on renewal and Markov processes, which allow simple modeling of stationary and non-stationary spike trains. Throughout the thesis, we consider mainly the Poisson processes with refractoriness (PPR), an adaptation of the Poisson processes, and the Gamma processes. In detail, we discuss their statistical properties. Furthermore, we present a typical form of discretization of time (used in SPADE) and point out the

significant implications of this approach.

We describe in the third chapter the particular complexity of spike trains obtained in electrophysiological recordings. As an example, we refer to the reach-to-grasp data set (Brochier et al. (2018)), which was already used for analyses with earlier versions of SPADE (Torre et al. (2016)).

SPADE is presented explicitly in the fourth chapter. We do not only provide an overview of the entire workflow and how spatio-temporal patterns are defined, but also explain in detail how the patterns are mined and statistically evaluated.

Chapter 5 presents the three dithering methods and compares them against one another. We contrast the standard approach of uniform dithering with a newly introduced adaptation, which takes refractoriness into account, and with the joint-ISI dithering. For uniform dithering, we analytically examine the generated surrogate spike trains in terms of their statistical properties. The same statistical properties are determined numerically for the other two methods, enabling us to compare the performance of the three methods.

In the last chapter, we develop a test case that allows a final assessment of the performances of the presented surrogate methods within SPADE, by evaluating them in terms of false positives and false negatives. We can compare the results obtained by the surrogate methods against a ground truth since we generate spike trains by employing the processes presented in the second chapter. Therefore, we discuss the reasons for the occurrences of false positives and false negatives. Consequently, depending on the firing rate, we recommend which dithering method to use for statistically robust pattern detection.

In conclusion, we summarize the results of this thesis and place them in the broader context of other surrogate methods, including non-dithering methods. Finally, we address the research questions raised by this thesis, and we give an outlook concerning upcoming studies.

Contents

1	Neuroscientific Background	7
1.1	Neurons and spikes	7
1.1.1	Spike trains	9
1.2	Neural Coding and cell assembly theory	9
1.2.1	Spatio-temporal Patterns	11
2	Point Processes	13
2.1	Definition of Point Process	13
2.2	Descriptive statistics	14
2.2.1	Firing rate	14
2.2.2	Autocorrelation and autocovariance	14
2.2.3	Spike count statistics	15
2.2.4	ISI-distribution	16
2.3	Spike Trains as finite-time point processes	17
2.4	Models of Point Processes	18
2.4.1	Renewal processes	18
2.4.2	Markov processes	20
2.4.3	Typical Spike Train models	22
2.4.3.1	Poisson process	22
2.4.3.2	Poisson process with refractoriness (PPR)	24
2.4.3.3	Gamma process	25
2.5	Discretization of Time	27
3	Electrophysiological Recordings	29
3.1	Reach-to-Grasp data set	29
4	SPADE	33
4.1	Definition of Spatio-Temporal Patterns	34
4.2	Mining of spatio-temporal patterns	34
4.3	P-Value spectrum and pattern spectrum filtering	36

5	Surrogate Methods	39
5.1	Uniform dithering	39
5.2	Uniform dithering with refractoriness	45
5.3	Joint-ISI dithering	46
5.4	Comparison of surrogate methods	48
5.4.1	Firing rate profile	49
5.4.2	Cross-correlation	49
5.4.3	ISI-distribution	53
5.4.4	Autocorrelation function	54
5.4.5	Clipped firing rate	54
5.4.6	Conclusion	58
6	False Positives and False Negatives	61
6.1	False positives	61
6.1.1	Ground truth null hypothesis	61
6.1.2	Null hypotheses using surrogate techniques	64
6.2	False negatives	70
6.2.1	Data generation with spatio-temporal patterns	70
6.2.2	Analysis of data with spatio-temporal patterns	71
7	Conclusion	77
7.1	Discussion	77
7.2	Implementation	81
7.3	Acknowledgements	82
A	Calculations for Point Processes	83
B	Calculations for Surrogate Methods	91
C	Miscellaneous to False Positives and False Negatives	101
	Bibliography	107
	List of Figures	113

Chapter 1

Neuroscientific Background

In this chapter, we briefly describe the processes behind the spiking activity of neurons. We pay special attention to why sequences of spikes are modeled using statistical point processes. Besides, we explain what spatio-temporal spiking patterns are and why hypotheses predict them.

It is not in the scope of this thesis to explain these topics in depth. In this respect, readers should refer to the cited neuroscience literature for more details.

1.1 Neurons and spikes

The mammalian brain consists mainly of two types of cells: neurons and glial cells. In the following, the focus will be limited to the consideration of neurons, since they are the principal cells responsible for sending electrical (and chemical) signals. The electrical signals are called action potentials and are explained in more detail later in this chapter.

A neuron consists of the three parts *soma*, *dendrites*, and an *axon* (including the axon terminals), as shown in figure 1.1. The soma is the cell body, which contains the nucleus. The axon leaves the soma and branches out at a certain point. The ends of these branches are called axon terminals, connected to dendrites of other cells via junctions called synapses (Gerstner et al. (2014)). The axon, with its terminal, is the part of the neuron responsible for propagating a signal from the soma to other neurons. Whereas, dendrites are extensions of a neuron that transmit electrochemical signals from other neurons to the soma. Thus, we can summarize the path of a neuron's signal as

$$\underbrace{\text{soma} \rightarrow \text{axon}}_{\text{presynaptic neuron}} \rightarrow \text{synapse} \rightarrow \underbrace{\text{dendrite} \rightarrow \text{soma}}_{\text{postsynaptic neuron}}.$$

For simplicity, we use the term neuron to refer to the soma and the term synapse for the entire path from one soma to another.

It is worth noting that there are different types of neurons. For our purpose, only the distinction between excitatory and inhibitory neurons is relevant. The name says it all: the signals from excitatory neurons have an amplifying effect on the membrane potential of the postsynaptic neuron, and those from inhibitory neurons have a weakening effect.




Figure 1.1: Schematic representation of a neuron showing the soma (cell body), the dendrites, and the axon with axon terminals. Figure from Salman et al. (2009).




Figure 1.2: Temporal evolution of membrane potential. Figure from Wikimedia Commons (2020).

The brain has a complex anatomy. For mammals, an essential part of it is the cerebral cortex, i.e., the outer (folded) layer of the neural tissue in the cerebrum. The cerebral cortex can be divided into cortical areas, which are thought to be primarily responsible for specific functions: e.g., the motor cortex for movement, the visual cortex for vision (Kandel et al. (2015)).

In the cerebral cortex, neurons are interconnected via synapses in such a way that they receive signals from and transmit signals to more than 10,000 neurons (Braitenberg and Schüz (1991)). Typically, the membrane potential describes the state of a neuron (Kandel et al. (2015)). The archetypal temporal dynamics of the membrane potential we show in figure 1.2.

There, it is shown the temporal evolution of the generation of a neuronal signal, by looking at the membrane potential of the neuron. The phenomenon starts with the membrane potential being at its resting state, called *resting potential*, which is around -70 mV. The membrane potential can increase due to electrical stimulation coming from other neurons. An increase in the membrane potential can have two possible consequences: it falls back to equilibrium, or it passes the *firing threshold*, which can vary from neuron to neuron, but it is usually around -55 mV. In the latter case, the neuron fastly depolarizes and then repolarizes. The depolarization and repolarization together are called an *action potential*. This action potential, with a temporal duration of 1 to 2 ms, is also called a *spike*. We say that the neuron *fires*. After the process of repolarization, the membrane potential takes on values smaller than the resting potential. The relaxation time back to the resting state is called the refractory period. We further differentiate the refractory period into two parts. In the first part, the absolute refractory period, no external input can elicit another action potential. While in the second part,

the relative refractory period, any stimulation will have a reduced effect (Kandel et al. (2015)).

In contrast to the figure 1.2, the membrane potential increases or decreases (depending on the input of excitatory or inhibitory neurons) by various incoming spikes whose combined effect can cause the membrane potential to exceed the threshold and the neuron to fire.

It is widely agreed that the spike's information lies only in its timing, neither in its amplitude nor in its shape (Kass et al. (2014)). Therefore, we have a particular interest in *spike trains*, sequences of spike times obtained from electrophysiological recordings. The process of extracting the spike times and assigning them to different neuron units is called *spike sorting*. This is still an error-prone process that is actively researched from many sides (Rey et al. (2015)¹) The methods, which we develop in this thesis, refer only to already spike-sorted data².

1.1.1 Spike trains

A *spike train* is defined as a sequence of ordered spike times $s = (t_1, \dots, t_N)$. In the usual experimental design, the recordings comprise different trials. In each trial is the spike train aligned to the presented external stimuli.

To measure the density of spikes over time, we define the firing rate as

$$\lambda(t) = \lim_{\epsilon \rightarrow \infty} \frac{1}{\epsilon} \text{Prob}(\text{spike in } [t, t + \epsilon]) \quad \text{for } t \geq 0, \epsilon > 0. \quad (1.1)$$

With a finite number of trials, this quantity cannot be evaluated exactly, but different estimation methods exist (Shinomoto (2010)). Other descriptive statistical properties we introduce in the next chapter. In order to model spike trains, stochastic point processes are used, which allow reproducing certain statistical observables.

Spike trains from simultaneous recordings of multiple neurons are called massively parallel spike trains. The algorithm SPADE, that will be introduced in chapter 4, was created to analyze this kind of experimental data. We denote k -th trial spike train of the n -th neuron will as s_n^k . In order to simplify the mathematical treatment, a spike train is written as a function of time, in the following manner:

$$s_n^k(t) = \sum_{i=1}^{N_n^k} \delta(t - t_i) \quad \text{with } t_i \in s_n^k \text{ and } N_n^k := \#\text{spikes in } s_n^k. \quad (1.2)$$

1.2 Neural Coding and cell assembly theory

How the brain processes information is still a controversial question in the neuroscientific community. There are two dominant ways to explain neural coding, i.e., the process of how neurons encode information (Brette (2015)).

¹There is also this recently finished master thesis in our institute: Sridhar (2020).

²In a certain way, one could estimate the impact of the errors due to spike sorting on the statistical analysis carried out in this thesis in further research using the analyses of Pazienti and Grün (2006).

The first one is called *rate coding*. In this framework, information is processed via firing rate modulation. Several studies showed the correlation between the dynamics of the firing rate and external stimuli or performed tasks (Georgopoulos et al. (1986); Lamme et al. (1998); Roelfsema et al. (2004)) supporting this theory.

The second one is called *temporal coding*. The basic idea is that the spike occurrences' precise timing plays an essential role in information processing. It was proposed by Abeles (1982) that neurons are *coincidence detectors*, i.e., a neuron is more likely to fire if it gets inputs synchronously from neurons connected by synapses than with asynchronous input. Experimental evidence in several studies that neurons work as coincidence detectors (Roy and Alloway (2001); Perez-Orive et al. (2004); Bender et al. (2006); Fino et al. (2010); Hong et al. (2012)) supports this hypothesis. The development of SPADE goes in the direction of finding further evidence for the hypothesis of temporal coding.

We introduce Hebbian Theory (Hebb (1949)) to illustrate why it is expected to find evidence for temporal coding in higher-order correlations. This theory provides a possible explanation for synaptic plasticity, i.e., the change in the weight of synapses over time: when one neuron fires and then a second connected neuron fires, the connection between these two neurons is strengthened. The strengthening of connections leads to the formation of cell assemblies, groups of strongly connected neurons, which are then able to reactivate in correspondence to a particular behavior.




Figure 1.3: Schematic representation of cell assemblies. Each triangle represents a neuron and each arrow a synapse. In each subfigure, one cell assembly is active (orange triangles). The synapses (blue arrows) propagate the spikes with different synaptic delays. As time progresses, other synapses (green arrows) activate other assemblies on the same space of neurons.




Figure 1.4: Spatio-temporal patterns in a raster plot of spike trains. Figure adapted from Quaglio et al. (2018).

1.2.1 Spatio-temporal Patterns

The presence of cell assemblies would mean that there are patterns of spiking activity extended in space and time: *spatio-temporal patterns* (STPs)³. The temporal extension of these patterns is expected due to the different synaptic delays. Nonetheless, before analyzing spatio-temporal patterns, neuroscientific research was focused on the analysis of unitary events (Grün et al. (2002a,b, 2010)), i.e., zero-lag correlations within a certain temporal precision. The spatio-temporal patterns dealt with in this thesis also allow for not zero-lag correlations in discretized time.

There are different methods for the detection of STPs on different time scales (Quaglio et al. (2018)). SPADE is a method for detecting exactly repeating patterns with a precision in the millisecond range.

The challenge of the statistical evaluation of the detected patterns in a temporal coding framework is equivalent to answer the question of how likely a pattern of a particular signature is to occur if the spike trains keep their statistics, while the correlations disappear.

³The question to what extent they are detectable in multi-unit recordings is currently being addressed in David Berling's master's thesis at our institute.

Chapter 2

Point Processes

In the previous chapter, we introduced spike trains as temporal sequences of action potentials of neurons. It is common to model spike trains by stochastic point processes. Therefore, this chapter gives an overview of stochastic point processes and defines statistical measures describing them. In particular, we discuss the Poisson process, an adapted version of it called the Poisson process with refractoriness (PPR), and the gamma process, whereby we distinguish between the stationary and non-stationary case. Finally, we evaluate two methods to discretize the time of the point process.

2.1 Definition of Point Process

A **stochastic point process** is an *ensemble of sets of time points* (Helias et al. (2015)):

$$\{t_1, \dots, t_i, \dots\} \text{ with } 0 \leq t_i \leq t_{i+1} \in \mathbb{R}^+. \quad (2.1)$$

In the given context, a set of time points corresponds to one realization of a spike train, where we call each time point a spike time. We will discuss the fact that spike trains are generally recorded on a finite time domain in section 2.3.

Each spike time t_i can be described by a *probability density function* that we denote as

$$\rho_i(t_i|t_{i-1}, \dots, t_1) \text{ with } \int_{t_{i-1}}^{\infty} dt_i \rho_i(t_i|t_{i-1}, \dots, t_1) = n_i \leq 1. \quad (2.2)$$

The case that the normalization $n_i < 1$ means that there is a probability of $1 - n_i$ that there is no further spike after t_{i-1} (van Vreeswijk (2010)). If not further specified, we will only treat processes for which $n_i = 1 \forall i$.

To derive statistical quantities we define an **average operator** over an ensemble of spike trains:

$$\langle f(\{t_i\}) \rangle_s := \int \prod_{j=1}^{\infty} (dt_j p_j(t_j|t_{j-1}, \dots, t_1)) f(\{t_i\}). \quad (2.3)$$

Therefore, to obtain an average value, we integrate over the p.d.f for each time point t_i .

As already mentioned in the previous chapter, for the analysis of time-dependent properties it is useful to represent a **spike train** as a *time-dependent function* (Helias et al. (2015)):

$$s(t) := \sum_{i=1}^{\infty} \delta(t - t_i) \text{ with } t \geq 0. \quad (2.4)$$

2.2 Descriptive statistics

Different statistical measures are commonly used in order to analyze the statistics of a spike train. In the following, we introduce all measures needed in this and the following chapters of the thesis.

2.2.1 Firing rate

The first measure, and one of the most fundamental, that we introduce is the **firing rate**. Generally, it is intuitively defined as the average number of spikes emitted by a neuron over time. Nonetheless, several similar (sometimes equivalent) definitions were given in literature when dealing mathematically with spike trains formalized as stochastic point processes (Dayan and Abbott (2001)). For example, we already defined the firing rate¹ as being proportional to the probability of having a spike in an infinitesimal interval (eq. 1.1). Another equivalent definition² *averages over the realizations* of the spike trains (Helias et al. (2015)):

$$\lambda(t) := \langle s(t) \rangle_s \text{ with } t \geq 0. \quad (2.5)$$

As described in the previous chapter, the firing rate is an important measure to describe spike trains' statistics.

The recorded neuronal firing rates vary in general over time from a biological perspective (Kass et al. (2014)). Nevertheless, we say a point process to be **stationary** when its rate does not change over time³:

$$\lambda(t) = \lambda > 0. \quad (2.6)$$

Stationary point processes are more straightforward to treat analytically and also give insights for the non-stationary ones.

2.2.2 Autocorrelation and autocovariance

The autocorrelation and autocovariance functions are used to calculate the impact of a spike at time t onto the probability of having a spike at time t' .

¹In the general theory of point processes, this quantity is called intensity function (Daley and Vere-Jones (1988)).

²We show the proof for the equivalence of the two firing rates definitions in the appendix section A.1.

³In the general theory of point processes, it is distinguished between first-order and second-order stationarity. The stationarity defined here would then correspond to the 1st order (Daley and Vere-Jones (1988)).

The **autocorrelation** function is defined as (Helias et al. (2015))

$$\psi(t, t') = \langle s(t)s(t') \rangle_s \text{ with } t, t' \geq 0. \quad (2.7)$$

Furthermore, it is useful to define the **autocovariance** function as

$$\begin{aligned} \tilde{\psi}(t, t') &:= \left\langle (s(t) - \langle s(t) \rangle_s)(s(t') - \langle s(t') \rangle_s) \right\rangle_s \text{ with } t, t' \geq 0 \\ &= \psi(t, t') - \lambda(t)\lambda(t'). \end{aligned} \quad (2.8)$$

2.2.3 Spike count statistics

Recorded spike trains are in general point processes on a finite time domain $[0, T]$.

On this *finite domain*, the **spike count** is defined as (Helias et al. (2015))

$$N(T) := \sum_{i=1}^{\infty} \Theta(T - t_i) \stackrel{(2.4)}{=} \int_0^T dt s(t) \quad (2.9)$$

With $\Theta(x)$ being the *Heaviside function* defined as $\Theta(x) = 1$ if $x \geq 0$ and 0 else. Intuitively, it represents the number of spikes emitted within the time interval $[0, T]$.

The **distribution of spike counts** is given by

$$p_N^T := \Pr(N = \#\text{spikes in } [0, T]) = \langle \delta_{N, N(T)} \rangle_s. \quad (2.10)$$

In the large time limit the distribution, of spike counts approximates a *Gaussian distribution* for most of the point processes (Onken et al. (2009); Rajdl and Lansky (2014)). Therefore, the distribution is then entirely described by the expectation value and the variance:

$$\mathbb{E}(N(T)) = \langle N(T) \rangle_s, \quad (2.11)$$

$$\text{Var}(N(T)) = \langle (N(T) - \langle N(T) \rangle_s)^2 \rangle_s. \quad (2.12)$$

These two values can also be represented as integrals over *the firing rate* respectively over the *autocovariance*:

$$\mathbb{E}(N(T)) \stackrel{(2.5)}{=} \int_0^T dt \lambda(t), \quad (2.13)$$

$$\text{Var}(N(T)) \stackrel{(2.8)}{=} \int_0^T dt \int_0^T dt' \tilde{\psi}(t, t'). \quad (2.14)$$

A typical measure for the *variability of the spike count* is the **Fano factor** (Nawrot et al. (2008)):

$$\text{FF}_T = \frac{\text{Var}(N(T))}{\mathbb{E}(N(T))}. \quad (2.15)$$

For spike trains originating from electrophysiological recordings, we replace the average operator $\langle \dots \rangle_s$ by the average over trials $\langle \dots \rangle_k$ introduced in section 2.3. In fact, repeating trials of an experiment are the sample counterpart of the different realizations of the same process.

2.2.4 ISI-distribution

In order to analyze the characteristics of the surrogate method we are going to introduce, we need to calculate statistics on the relationship between successive spikes. An **inter-spike interval** (ISI) is defined as the time interval given by two subsequent spikes. We can thus define the **inter-spike interval distribution**:

$$f(\tau) = \lim_{T \rightarrow \infty} c_1(T) \left\langle \sum_{i=1}^{\infty} \delta(\tau - (t_{i+1} - t_i)) \Theta(T - t_{i+1}) \right\rangle_s \quad (2.16)$$

$$\text{with } c_1(T) = \frac{1}{E(N(T)) - \langle \Theta(T - t_1) \rangle_s} \approx \frac{1}{E(N(T)) - 1} \quad (2.17)$$

Our notation for the average over the ISI distribution is $\langle \dots \rangle_\tau = \int_0^\infty d\tau f(\tau) \dots$. The mean ISI and the variance are

$$\mu_\tau := \langle \tau \rangle_\tau \quad (\sigma_\tau)^2 := \langle \tau^2 \rangle_\tau - (\langle \tau \rangle_\tau)^2. \quad (2.18)$$

A commonly used indicator of the regularity of a spike train is the **coefficient of variation**, defined as (van Vreeswijk (2010))

$$C_V := \frac{\sigma_\tau}{\mu_\tau}. \quad (2.19)$$

Additionally, we define the **joint-ISI distribution**, which is particularly interesting for us, since one of the surrogate methods taken into observation will be the joint-ISI dithering (introduced in section 5.3). This distribution is a two-dimensional probability density that gives the joint probability to have an inter-spike interval τ followed subsequently by an interval τ' .

$$f(\tau, \tau') = \lim_{T \rightarrow \infty} c_2(T) \left\langle \sum_{i=1}^{\infty} \delta(\tau - (t_{i+1} - t_i)) \delta(\tau' - (t_{i+2} - t_{i+1})) \Theta(T - t_{i+2}) \right\rangle_s \quad (2.20)$$

$$\text{with } c_2(T) = \frac{1}{E(N(T)) - \langle \Theta(T - t_1) + \Theta(T - t_2) \rangle_s} \approx \frac{1}{E(N(T)) - 2} \quad (2.21)$$

Since it has been observed that the coefficient of variation is not an appropriate measure for regularity when the process is non-stationary (van Vreeswijk (2010)), we propose another more robust quantity, called **CV2**:⁴

$$C_{V2} = 2 \left\langle \frac{|\tau' - \tau|}{\tau' + \tau} \right\rangle_{\tau, \tau'}. \quad (2.22)$$

⁴The average operator of the joint-ISI distribution is $\langle \dots \rangle_{\tau, \tau'} = \int d\tau \int d\tau' f(\tau, \tau') \dots$

2.3 Spike Trains as finite-time point processes

As we have seen, electrophysiological recordings comprise a finite number of trials performed for a stimulus or task. The spike train of each trial is thus recorded in a time window $[0, T]$. Statistical quantities like the firing rate or the ISI distribution are correspondingly estimated either as average over trials or from a single trial. Trials are indicated throughout the chapter, when necessary, with the index k .

Since the firing rate is, in general, a smooth, continuous function, it is determined using kernel density estimation over the spike train (Shinomoto (2010)).

On a trial-by-trial basis, the firing rate would be (Shinomoto (2010))

$$\lambda^k(t) = (s^k \otimes h)(t). \quad (2.23)$$

The average over trials $\langle \dots \rangle_k$ yield (Shinomoto (2010))

$$\lambda(t) = (\langle s^k \rangle_k \otimes h)(t). \quad (2.24)$$

In principle, several normalized kernels $h(t)$ can be used (Shinomoto (2010)). We decided to use only a Gaussian function $\mathcal{N}(t, 0, \sigma^2)$ here, with suitably chosen standard deviation σ . Note that to choose the parameter σ is more challenging in the single-trial condition (Cunningham et al. (2009)).

To estimate the ISI distribution and the joint-ISI distribution using a kernel density estimation technique is more subtle than for the firing rate since these distributions are defined on the positive domain of real space. Hence, problems arise in estimating the range of small intervals τ of these distributions.

Therefore, these distributions are estimated using histograms, where it is then necessary to choose a suitable discretization interval, or bin size $b > 0$:

$$f_n = c_1 \int_{nb}^{(n+1)b} d\tau \left\langle \sum_{i=1}^{N^k-1} \delta(\tau - (t_{i+1} - t_i)) \right\rangle_k \quad (2.25)$$

with $c_1 = \frac{1}{\langle N^k \rangle_k - \langle \Theta(N^k - 1) \rangle_k},$

$$f_{n,n'} = c_2 \int_{nb}^{(n+1)b} d\tau \int_{n'b}^{(n'+1)b} d\tau' \left\langle \sum_{i=1}^{N^k-2} \delta(\tau' - (t_{i+2} - t_{i+1})) \delta(\tau - (t_{i+1} - t_i)) \right\rangle_k \quad (2.26)$$

with $c_2 = \frac{1}{\langle N^k \rangle_k - \langle \Theta(N^k - 2) + \Theta(N^k - 1) \rangle_k}.$

and with N^k indicating the spike count of the k -th trial.

Nevertheless, in section 5.3, it will be necessary also to estimate the joint-ISI distribution for spike trains with a small number of spike counts. For this purpose, we introduce Gaussian filtering with a narrow kernel width σ as:

$$f_{n,n'}^\sigma \propto \sum_{m,m'} \mathcal{N}(nb, mb, \sigma^2) \mathcal{N}(n'b, m'b, \sigma^2) f_{m,m'} \text{ for } n, n' \geq n^*. \quad (2.27)$$

In order not to change the minimal inter-spike interval⁵ and using an initial guess of the refractory period τ_r^{init} , we set the minimal bin position n^* as

$$n^* = \min \left(\min_n (f_n > 0), \left\lfloor \frac{\tau_r^{\text{init}}}{b} \right\rfloor \right). \quad (2.28)$$

2.4 Models of Point Processes

In the following, we present the point process models of renewal processes and Markov processes. The former is an extreme case of the latter. Renewal processes can be well used to model stationary spike train data and Markov processes for non-stationary spike trains. We will cover all relevant features like the ISI-distribution and the autocorrelation function.

Then, we introduce typical renewal (and Markov) processes with a particular focus on the Poisson process. Next, we concentrate on the Poisson process with refractoriness (PPR) and gamma process. Both return to the Poisson process for appropriately chosen parameters. These spike train models we use afterward to generate test data for the analysis of SPADE.

2.4.1 Renewal processes

A **renewal process** is a process where the probability density of the spike time t_i depends only on the interval to the previous spike time t_{i-1} . In mathematical terms, this is expressed as⁶

$$\begin{aligned} \rho_i(t_i | t_{i-1}, \dots, t_1) &= p(t_i - t_{i-1}) \quad \forall i > 1 \\ \text{with } \int_0^\infty d\tau p(\tau) &= 1 \text{ and } p(\tau) = 0 \text{ for } \tau < 0 \end{aligned} \quad (2.29)$$

The probability density function $p(\tau)$ is called the **renewal inter-spike interval (ISI) distribution**. It is easy to prove that this corresponds to the general definition of the ISI distribution from eq. (2.16): $p(\tau) = f(\tau)$. Thus, it is possible to determine the density

⁵This will later be necessary to conserve the refractory period.

⁶This definition is less strict than commonly used like in Daley and Vere-Jones (1988).

of the i -th spike time by recursively integrating the dependencies of prior spikes:

$$\rho_i(t_i) = \int dt_{i-1} dt_{i-2} \dots dt_2 dt_1 p(t_i - t_{i-1}) p(t_{i-1} - t_{i-2}) \dots p(t_2 - t_1) \rho_1(t_1) \quad (2.30)$$

$$= \begin{cases} \rho_1(t_1) & \text{if } i = 1 \\ (\rho_1 \circledast p^{(i-1)\circledast})(t_i) & \text{if } i > 1 \end{cases}$$

with $p^{(i)\circledast}(\tau) = \underbrace{(p \circledast \dots \circledast p)}_{i \text{ times}}(\tau)$ for $i \in \mathbb{N}$ (2.31)

Since, by definition, here each spike depends uniquely from its previous, it follows that subsequent ISIs are independent. For this reason, the joint-ISI distribution factorizes:

$$f(\tau, \tau') = f(\tau) \cdot f(\tau') = p(\tau) \cdot p(\tau'). \quad (2.32)$$

We define the **positive- τ conditional firing rate** as

$$\lambda_c^+(\tau) := \sum_{i=1}^{\infty} p^{(i)\circledast}(\tau). \quad (2.33)$$

This quantity indicates that for a given spike, the probability of having a spike in the interval $[\tau, \tau + d\tau]$ after that spike is $\lambda_c^+(\tau)d\tau$.

In agreement with the positive- τ conditional firing rate, we define the conditional firing rate as

$$\lambda_c(\tau) := \lambda_c^+(\tau) + \lambda_c^+(-\tau), \quad (2.34)$$

which thus also indicates the probability of having a spike before a given spike.

We can derive a relationship between firing rate and positive- τ conditional firing rate as (see appendix eq: A.1)

$$\lambda(t) = \rho_1(t) + (\rho_1 \circledast \lambda_c^+)(t). \quad (2.35)$$

The literature distinguishes between ordinary and equilibrium renewal processes (van Vreeswijk (2010)). An ordinary renewal process is a process that starts with an uncounted spike at $t = 0$. So that for the first spike holds $\rho_1(t_1) = p(t_1)$. While an equilibrium renewal process is a process being from time zero already in its equilibrium state, or differently formulated: the first spike is distributed such that the process is stationary (Asmussen (2003)). Analytically:⁷

$$(\rho_1 \circledast \lambda_c^+)(t) + \rho_1(t) = \lambda \quad (2.36)$$

⁷The equilibrium renewal process is usually defined in such a way that there is a spike at $t \ll 0$, but the time window under consideration only starts at $t = 0$ (van Vreeswijk (2010)). However, our definition is easier to use from an analytical point of view and is equivalent except for extreme cases.

In this manuscript, we will only use the equilibrium renewal process since the property of stationarity can be quite advantageous.

Using the stationarity and the formula for the firing rate in renewal processes (see eq. (2.35)), as derived in the appendix in eq. (A.8), the density of the first spike becomes

$$\rho_1(t_1) = \lambda \Theta(t_1) \left(1 - \int_0^{t_1} d\tau p(\tau) \right) \quad (2.37)$$

This density of the first spike will later help us to describe the effect of discretization of time (see section 2.5).

The autocorrelation function and the autocovariance simplify to (see appendix eq. (A.2))

$$\psi(t, t') = \lambda \delta(t - t') + \lambda \lambda_c(t - t') \quad (2.38)$$

$$\tilde{\psi}(t, t') = \lambda \delta(t - t') + \lambda(\lambda_c(t - t') - \lambda) \quad (2.39)$$

Therefore, it is sufficient to analyze the conditional firing rate in order to investigate the auto-structure of the renewal processes.

Another well-known fact for renewal processes is that the Fano factor is related to the C_V by (Cox and Lewis (1966))

$$\lim_{T \rightarrow \infty} \text{FF}_T = (C_V)^2 \quad (2.40)$$

Thus, the relationship between the two is used to test if spike trains recorded in electrophysiological recordings are underlined by renewal processes (Nawrot et al. (2008)). If the deviation is too large, it is not justified to use renewal processes when dealing with stationary spike trains.

2.4.2 Markov processes

A Markov process is a process where the probability density of the spike time t_i depends only on the previous spike time t_{i-1} . Thus, the renewal process defined above is a particular case of a Markov process.

The probability density of a spike time t_i becomes then

$$\rho_i(t_i | t_{i-1}, \dots, t_1) = p(t_i | t_{i-1}) \quad \forall i > 0. \quad (2.41)$$

$$\text{with } \int_0^\infty dt p(t | t') = 1 \text{ and } p(t | t') = p(t | t') \cdot \Theta(t - t') \quad (2.42)$$

We call this quantity $p(t_i | t_{i-1})$ the **Markov density function**.

The ISI-distribution and the joint-ISI distribution from eqs. (2.16 & 2.20) become following to appendix (see sec. A.3)

$$f(\tau) = \lim_{T \rightarrow \infty} \frac{\int_0^T dt \lambda(t) p(t + \tau | t)}{\int_0^T dt \lambda(t)} \quad (2.43)$$

$$f(\tau, \tau') = \lim_{T \rightarrow \infty} \frac{\int_0^T dt \lambda(t) p(t + (\tau + \tau') | t + \tau) p(t + \tau | t)}{\int_0^T dt \lambda(t)} \quad (2.44)$$

We will see later that these quantities are difficult to interpret since we did not impose any constraint on the firing rate. The distributions can still be best understood by approximating the firing rate as locally stationary.

Consider the case that the firing rate consists of locally stationary firing rates λ_i on time intervals ΔT_i . On each of these intervals, the Markov process reduces to a renewal process with renewal inter-spike interval function $p(\tau; \lambda_i)$. When neglecting the ISIs that will result from the boundaries between different stationary stages, the ISI and joint-ISI distribution reduce to

$$f(\tau) = \frac{\sum_i \lambda_i \Delta T_i p(\tau; \lambda_i)}{\sum_i \lambda_i \Delta T_i}, \quad (2.45)$$

$$f(\tau, \tau') = \frac{\sum_i \lambda_i \Delta T_i p(\tau; \lambda_i) p(\tau'; \lambda_i)}{\sum_i \lambda_i \Delta T_i}. \quad (2.46)$$

with $p(\tau; \lambda_i) = p(t + \tau | t)$ for t in ΔT_i

This approximation is only valid, if $\Delta T_i \gg \frac{1}{\lambda_i} \forall i$.

It can be derived, similarly as for the case of renewal processes, that the probability distribution of the first spike is:

$$\rho_1(t_1) = \lambda(t_1) - \int_0^{t_1} dt p(t_1 | t) \lambda(t) \quad (2.47)$$

Similar to eq. (2.30), the distribution of the i -th spike becomes

$$\rho_i(t_i) = \int_0^{t_i} dt_{i-1} p(t_i | t_{i-1}) \rho_{i-1}(t_{i-1}) \quad \forall i > 1 \quad (2.48)$$

Thus, the **firing rate** can be expressed as

$$\lambda(t) = \sum_{i=1}^{\infty} \rho_i(t) \quad (2.49)$$

The **autocorrelation** and **autocovariance** become

$$\psi(t, t') = \lambda(t)\delta(t - t') + \lambda(t')\lambda_c(t|t') \quad (2.50)$$

$$\tilde{\psi}(t, t') = \lambda(t)\delta(t - t') + \lambda(t')(\lambda_c(t|t') - \lambda(t)) \quad (2.51)$$

Finally, the **conditional firing rate** for a Markov process boils down to being the firing rate at time t in the presence of a spike at time t' :

$$\lambda_c(t|t') := \frac{\lambda_c^+(t|t')\lambda(t') + \lambda_c^+(t'|t)\lambda(t)}{\lambda(t')} \quad (2.52)$$

$$\text{with } \lambda_c^+(t|t') := \sum_{i=1}^{\infty} p_i(t|t') \quad (2.53)$$

$$\text{and } p_1(t|t') = p(t|t') \text{ and } p_i(t|t') = \int_{t'}^t ds p(t|s)p_{i-1}(s|t') \forall i > 1$$

We have summarized the information on Markov processes here mainly for the sake of completeness. This information will become especially useful when later research extends the test case of false positives and false negatives to non-stationary spike trains.

2.4.3 Typical Spike Train models

We introduce here three point process models, which have a stationary version that is renewal and a non-stationary version that is a Markov process. The Poisson and gamma processes are widely studied in literature and standard approaches for modeling spike trains (Nawrot et al. (2008)). The Poisson Process with refractoriness (PPR) is less analyzed⁸ but will be of great importance for our purposes next.

In the following, we will limit ourselves to summarize the most relevant properties, "table-like".

2.4.3.1 Poisson process

The Poisson Process is the most studied and simplest point process.

Stationary

The stationary version, which is a renewal process, can be defined in several ways. Within the current framework, we define the Poisson Process via its ISI distribution, which is an exponential decay (van Vreeswijk (2010)):

$$p(\tau) = \lambda \exp(-\lambda\tau)\Theta(\tau). \quad (2.54)$$

We show this ISI distribution in figure 2.1, where it is denoted as a gamma process (sec. 2.4.3.3) with shape factor $\gamma = 1$.

⁸Among the articles published about PPR spike trains, Deger et al. (2012) is particularly noteworthy.

Another equivalent definition refers to the spike count distribution. The spike count distribution corresponds for every interval $[s_1, s_2]$ to a Poisson distribution (Daley and Vere-Jones (1988)):

$$p_N^{[s_1, s_2]} := \Pr(N = \#\text{spikes in } [s_1, s_2]) = \frac{(\lambda \cdot (s_2 - s_1))^N}{N!} \exp(-\lambda \cdot (s_2 - s_1)). \quad (2.55)$$

Additionally, the spike count distributions for two non-overlapping intervals are independent.

The distribution of the first spike is expressed as (van Vreeswijk (2010))

$$\rho_1(t_1) = p(t_1). \quad (2.56)$$

In correspondence to section 2.4.1, we have an equivalence between ordinary and equilibrium renewal processes in the case of the Poisson process. This follows from what is called the *memoryless property* of the process, that is reflected in its autocovariance:

$$\tilde{\psi}(t, t') = \lambda \delta(t - t'). \quad (2.57)$$

Thus, the presence of a spike at time t does not change the probability of having a spike at $t' \neq t$.

Moreover, all measures of regularity introduced before are constant and equal to 1 (van Vreeswijk (2010)) for a homogeneous Poisson process:

$$C_V = C_{V2} = \text{FF}_T = 1. \quad (2.58)$$

Consequently, it is a drawback of modeling a spike train as a Poisson process that the regularity measures (C_V , C_{V2} , FF) cannot be fit to those of the data.

Non-stationary

An extension of a stationary Poisson process, including a non-stationary firing rate, is possible as a Markov process. Therefore, some replacements need to be applied: $\lambda\tau$ is replaced by $\int_{t_{i-1}}^{t_i} dt \lambda(t)$ and λ by $\lambda(t_i)$. Thus, the Markov density function becomes

$$p(t_i | t_{i-1}) = \lambda(t_i) \exp \left(- \int_{t_{i-1}}^{t_i} dt \lambda(t) \right) \Theta(t_i - t_{i-1}). \quad (2.59)$$

The integral over the firing rate is also referred to as operational time $t' = \int_0^t dt \lambda(t)$ (Cox and Isham (1980); Nawrot et al. (2008); Louis et al. (2010b)). If the time of the non-stationary Poisson process is transformed into this operational time, in this operational time, the process will be stationary and renewal. Therefore, it is also called a rate-modulated renewal process (Nawrot et al. (2008)).

2.4.3.2 Poisson process with refractoriness (PPR)

The Poisson process with refractoriness (PPR) is an adaptation of the Poisson process, which includes refractoriness. In the context of this master thesis, this process was intensively dealt with and implemented in the software package elephant⁹ for the stationary and non-stationary case.

Stationary

In order to take into account an absolute **refractory period** $\tau_r > 0$, the exponential decay of the ISI distribution of the Poisson Process is shifted by this τ_r . Consequently, an **effective firing rate** $\tilde{\lambda} > \lambda$ needs to be introduced to keep the expected number of spikes. This effective firing rate corrects for the time windows after each spike, in which no spike can be emitted. Given that, we have λT spikes on average, the “blocked” time windows sum up to $\lambda T \tau_r$. The average spike count should correspond to the effective firing rate acting in the shortened time window $T - \lambda T \tau_r$ and thus:

$$\tilde{\lambda}(T - \lambda T \tau_r) = \lambda T \Leftrightarrow \tilde{\lambda} = \frac{\lambda}{1 - \lambda \tau_r}. \quad (2.60)$$

The **ISI distribution** follows as

$$p(\tau) = \tilde{\lambda} \exp(-\tilde{\lambda} \cdot (\tau - \tau_r)) \Theta(\tau - \tau_r). \quad (2.61)$$

We show this ISI distribution in figure 2.1 for various refractory periods with a fixed firing rate.

The distribution for the first spike results from eq. (2.37):

$$\rho_1(t_1) = \begin{cases} \lambda & \text{for } 0 \leq t_1 \leq \tau_r \\ \lambda \exp(-\tilde{\lambda} \cdot (t_1 - \tau_r)) & \text{for } t_1 > \tau_r \\ 0 & \text{for } t < 0 \end{cases}. \quad (2.62)$$

The C_V (from Deger et al. (2012)) and C_{V2} (appendix eqs. (A.10 & A.11)) are given as

$$C_V = \frac{\lambda}{\tilde{\lambda}} = 1 - \lambda \tau_r, \quad (2.63)$$

$$C_{V2} = \int_{2\tau_r}^{\infty} dt \tilde{\lambda}^2 \frac{(t - 2\tau_r)^2}{t} \exp(-\tilde{\lambda}(t - 2\tau_r)) = 1 - 2\lambda \tau_r + \mathcal{O}((\tau_r)^2). \quad (2.64)$$

It is noteworthy that both the C_V and C_{V2} depend on the refractory period τ_r .

A closed formula for the autocorrelation and the Fano factor also can be found in Deger et al. (2012).

⁹Elephant is an open-source Python software package to analyze electrophysiological recordings. It is mainly developed in our institut. <https://elephant.readthedocs.io>.




Figure 2.1: ISI distributions for PPR and gamma processes using a firing rate of $\lambda = 100$ Hz. Note that the gamma process with $\gamma = 1$ is the Poisson process.

Non-stationary

The replacement needed to extend the stationary PPR to non-stationarity is not as straightforward as for the Poisson process.

In accordance with equation (2.60), the **effective firing rate** needs to satisfy:

$$\int_0^T dt \tilde{\lambda}(t) \left(1 - \int_{t-\tau_r}^t ds \lambda(s) \right) = \int_0^T dt \lambda(t) \Rightarrow \tilde{\lambda}(t) = \frac{\lambda(t)}{1 - \int_{t-\tau_r}^t ds \lambda(s)}. \quad (2.65)$$

Hence, the effective firing rate at time t depends on the firing rate in the interval $[t - \tau_r, t]$. Further substituting $\tilde{\lambda} \cdot (\tau - \tau_r) \rightarrow \int_{t_{i-1}+\tau_r}^{t_i} dt \tilde{\lambda}(t)$ and $\tilde{\lambda} \rightarrow \tilde{\lambda}(t_1)$ yields the **Markov density function**:

$$p(t_i | t_{i-1}) = \tilde{\lambda}(t_i) \exp \left(- \int_{t_{i-1}+\tau_r}^{t_i} dt \tilde{\lambda}(t) \right) \Theta(t_i - (t_{i-1} + \tau_r)). \quad (2.66)$$

Thus, this process is not a rate-modulated renewal process, since it does not depend on an interval in operational time.

2.4.3.3 Gamma process

Gamma processes can be used to model both regular and bursty spike trains. In the regular case, the inter-spike intervals get closer to the mean ISI $\frac{1}{\lambda}$, while for the bursty

case, it is more likely that the intervals are much smaller or higher.

Stationary

We define the *gamma process* via its **ISI distribution** that includes the shape factor $\gamma > 0$ (van Vreeswijk (2010)):

$$p^\gamma(\tau) = \gamma \lambda \frac{(\gamma \lambda \tau)^{(\gamma-1)}}{\Gamma(\gamma)} \exp(-\gamma \lambda \tau) \Theta(\tau). \quad (2.67)$$

This ISI distribution can be seen in figure 2.1 for various shape factors with a fixed firing rate.

With the shape factor, it is possible to distinguish between regular processes $\gamma > 1$ and bursty ones $\gamma < 1$. If $\gamma = 1$ it recovers the Poisson process.

For the p.d.f of the first spike (eq. (2.37)), it is necessary to introduce the upper and lower incomplete gamma functions as

$$\Gamma(\gamma, x) := \int_x^\infty ds s^{\gamma-1} e^{-s} \quad \bar{\Gamma}(\gamma, x) := \int_x^\infty ds s^{\gamma-1} e^{-s}. \quad (2.68)$$

Thus,

$$\rho_1^\gamma(t_1) = \lambda \frac{\Gamma(\gamma, \gamma \lambda t_1)}{\Gamma(\gamma)} \Theta(t_1). \quad (2.69)$$

The **measures of regularity** for the stationary gamma process are (van Vreeswijk (2010)):

$$C_V = \frac{1}{\sqrt{\gamma}}, \quad (2.70)$$

$$C_{V2} = \frac{\Gamma(2\gamma)}{\gamma (2^{\gamma-1} \Gamma(\gamma))^2}. \quad (2.71)$$

An analysis of the Fano factor for stationary gamma processes can be found in Rajdl and Lansky (2014).

Non-stationary

Here the same replacement is applied as for the Poisson process, leading the **Markov density function** to be

$$p^\gamma(t|t') = \gamma \lambda(t) \frac{\left(\gamma \int_{t'}^t ds \lambda(s) \right)^{(\gamma-1)}}{\Gamma(\gamma)} \exp \left(-\gamma \int_{t'}^t ds \lambda(s) \right) \Theta(t - t'). \quad (2.72)$$

Thus, depending on an interval in operational time $\int_{t'}^t ds \lambda(s)$, the non-stationary gamma process can be also classified as a rate-modulated renewal process.

2.5 Discretization of Time

We introduce discretization methods for spike trains while keeping in mind the question of the temporal precision of the neuronal code and the spatio-temporal patterns. There are two standard models of discretization, which are called **binning** and **clipping**. A spike train, typically modeled as a continuous-time process, is discretized by dividing the interval of time where it is defined into contiguous sub-intervals called bins of equal length (bin size). The corresponding discretized spike train can contain in each bin either the number of spikes (binning) or a 1 if the bin is occupied (clipping). Of course, a 0 is assigned for both cases if the bin is empty. The clipping has the numerical advantage that spike trains are converted into binary data. However, it is likely that not all spikes are counted. Thus, the firing rate is reduced. This effect is small with suitably selected bin size due to the refractory period in biological spike trains but can become significant when using surrogate methods.

Since the SPADE algorithm uses clipped spike trains (see chap. 4), it will be essential to pay attention to this concerning the surrogate methods. To give a first insight, we analyze here the **clipped firing rates** for renewal point processes.

Using the bin size $b > 0$, the spike train is **clipped in this way** that in the m -th bin it becomes

$$s_m^{\text{clip}} = \begin{cases} 1 & \text{if } \int_{mb}^{(m+1)b} dt s(t) \geq 1 \\ 0 & \text{else} \end{cases}. \quad (2.73)$$

For renewal processes, we can calculate the clipped firing rate by integrating one bin over the first spike's probability density, since the integral indicates the probability of having at least one spike in this first bin. Furthermore, since we chose the processes to be in equilibrium and stationary, the **clipped firing rate** is *constant* for all bins.

$$\lambda^{\text{clip}} = \frac{1}{b} \langle s^{\text{clip}} \rangle = \frac{1}{b} \int_0^b dt \rho_1(t) dt \quad (2.74)$$

This results in the following **clipped firing rates** for the three point process models (see appendix A.5):

$$\lambda_{\text{Pois}}^{\text{clip}} = \frac{1}{b} (1 - \exp(-\lambda b)), \quad (2.75)$$

$$\lambda_{\text{PPR}}^{\text{clip}} = \begin{cases} \lambda & \text{for } b \leq \tau_r \\ \lambda \frac{\tau_r}{b} \exp(-\tilde{\lambda}(b - \tau_r)) + \frac{1}{b} (1 - \exp(-\tilde{\lambda}(b - \tau_r))) & \text{for } b > \tau_r \end{cases}, \quad (2.76)$$

$$\lambda_{\gamma}^{\text{clip}} = \lambda \frac{\Gamma(\gamma, \gamma \lambda b)}{\Gamma(\gamma)} + \frac{1}{b} \left(\frac{\bar{\Gamma}(\gamma + 1, \gamma \lambda b)}{\Gamma(\gamma + 1)} \right). \quad (2.77)$$

Here, we make use of the incomplete gamma functions defined in equation (2.68).




Figure 2.2: Clipped firing rate for PPR (left) and gamma processes (right) using a bin size of $b = 3$ ms. For the different parameters of the spike trains, we use different colors. Note that the gamma process with $\gamma = 1$ is the Poisson process.

As can be seen in Figure 2.2, the effect of clipping is much stronger for gamma processes than for PPR. Furthermore, the percentage of the firing rate decreases more for higher firing rates. This can be understood by the fact that higher firing rates compress the ISI distributions, making it more likely to get inter-spike intervals smaller than the bin size. The bin size of 3 ms has been used as a standard parameter for SPADE analyses in recent years, as the scale of temporal precision was assumed to be in this range (Torre et al. (2013, 2016)). In possible analyses of the optimal bin size, special attention should be paid to the effects on the clipped firing rate, which further decreases the larger the bin size.

For the non-stationary Markov processes, one can approximate them well to be locally stationary within each bin since the variation of firing rate is negligible in the range of a few milliseconds. Therefore, it is also possible to estimate the clipped firing rate for non-stationary processes. In further research, this is needed when estimating the clipped firing rate for data from electrophysiological recordings.

Chapter 3

Electrophysiological Recordings

The technique of recording the electrical activity of neurons is called electrophysiology. These recordings are the experimental basis for the understanding of the mammalian brain. In order to analyze the relation between neural activity and behavior, *in vivo* recordings are needed. These are obtained by experiments using multielectrode arrays implanted in the cerebral cortex. We call massively parallel spike trains, the spike trains extracted from these recordings. The SPADE-method was developed to detect spatio-temporal in this type of spike trains.

Any statistical analysis is designed to address and fit the statistical complexity of the observed data. This is particularly important for electrophysiological recordings, where the data is highly variable in time and multidimensional. To illustrate this, we introduce an example data set already analyzed with a previous version of SPADE (Torre et al. (2016)). We also display its most relevant statistical features, such as firing rate and regularity measures.

3.1 Reach-to-Grasp data set

The data set that we call reach-to-grasp data set was recorded from Thomas Brochier and his collaborators at the Institute of Neuroscience de la Timone (Marseille) in the motor cortex of two macaque monkeys (Brochier et al. (2018)). The cited publication does not only describe the experiment but also makes accessible a part of the recordings to the scientific community. The monkeys had to perform an instructed delayed reaching and grasping of an object, while an implanted 10x10 Utah electrode array was recording the cerebral activity in the motor cortex. This recording did not only serve to extract massively parallel spike trains but also yield recordings of the local field potential.

In each of the trials, the monkey had to perform the reach-to-grasp task in one of four possible ways, indicated to the monkey by visual cues. The recordings are organized into sessions of about 10 minutes. The trials within a session are typically divided into epochs of 500 ms, related to specific behavioral contexts (Torre et al. (2016)).

According to the findings in Torre et al. (2016), we only show the data of the epoch in which the movement happens. First of all, it is interesting to have a look at the firing

rate profiles. We can see in figure 3.1 that the firing rates are notably non-stationary, and they strongly vary across neurons. To see how the firing rates correlate to different behavioral contexts or how the dynamics are for an entire session, readers may refer to Riehle et al. (2018).

The fact that firing rates vary across neurons is critical for the detection and the statistical evaluation of spatio-temporal patterns and thus for the design of the surrogate method. In order to understand this, it needs to be considered that the expected occurrences of by-chance patterns are proportional to the product of firing rates (Pipa et al. (2013)¹). SPADE has shown of being able to deal with heterogeneous firing rates for the case of simulated data (Quaglio et al. (2017); Stella et al. (2019)). The last two studies take into consideration firing rates up to 25 Hz. However, in the reach-to-grasp data set, we observe an inevitable part of neuronal units exceeding this strongly.

It also needs to be taken into account that the firing rate of a single neuron unit is not constant over time, which does not directly translate to be a result of rate coding (which we had discussed in sec. 1.2). However, an indication for the rate coding hypothesis would be the correlation between the temporal dynamics of the neuron-averaged mean of the firing rates for the whole trial and the corresponding behavioral context, which can be seen in Riehle et al. (2018).

Since our purpose is to detect correlation not emerging uniquely by the firing rate, a clear statistical distinction is necessary between by-chance patterns as a result of the firing rate co-modulations and significant patterns of rate co-modulated neuron units. This distinction will be made by designing an appropriate null hypothesis for the statistical test (see chaps: 4 & 5).

In order to further analyze the data set introduced here, we take into account the statistical measures of the previous chapter (sec: 2.2). Specifically, we inspect C_V , C_{V2} , FF , and the minimal ISI for each neuron within the movement epoch for all trials of the same sessions analyzed previously. The coefficient of variation C_V is the first indicator of the regularity of spike trains (see fig: 3.2). Since by definition it will be influenced by the dynamics of the firing rates - that vary strongly - we also show the C_{V2} . It is noteworthy that for both monkeys for average firing rates over 20 Hz, the C_{V2} is smaller than 1. Thus, it can be summarized that all “active” neuron units are regular. A similar result can be seen in Riehle et al. (2018), where it is shown that the C_{V2} is usually smaller during movement than when the monkey is waiting.

The Fano factor, shown in the lower-left corner of the figure, interestingly seems to be on average lower values for monkey 2 than for monkey 1. Since the spike trains cannot be modeled by renewal processes - as a consequence of their non-stationarity - we do not further analyze the relationship of C_V and FF (see eq. 2.40).

The last panel of the two figures (fig. 3.2) shows the minimal inter-spike interval (ISI) in the data. We analyze this quantity in addition to the C_V and the C_{V2} to see in how far it is reasonable to model this spike trains with PPR spike trains instead of

¹In that paper the relation of the expected occurrences to the firing rate is only shown for zero-lag patterns. However, starting from their reasonings it is easy to show this for also for temporally extended patterns.




Figure 3.1: Firing rate profile of neurons of two sessions of the reach-to-grasp experiment of two different monkeys (left and right panel, session *i140703* – 001 for monkey N and session *l101210* – 001 for monkey L). An epoch of 500 ms is shown, aligned to the switch release (SR), which is the time point the movement starts. The firing rate is estimated as described in sec. 2.3 using a kernel width of $\sigma = 25$ ms and averaging over trials. For both monkeys, we highlight the firing rate profiles for the neurons with the five highest time-averaged firing rates in colors. The other rate profiles are shown in grey. On the left side, we see that most of the firing rates are below 25 Hz, and only a few neuron units exceed this limit. Only one neuron unit exceeds an instantaneous firing rate of 85 Hz; this one goes up until 175 Hz. On the right side, no neuron unit exceeds an instantaneous firing rate of 90 Hz.

what is typically done in the literature with gamma spike trains (Nawrot et al. (2008); Pipa et al. (2013)).

For most neurons, the minimal ISI is close to the value of the spike wavelength used for the spike sorting (Brochier et al. (2018)), which are 1.3 ms and 1.6 ms, depending on the monkey. For these neurons, it is thus reasonable to argue that their minimal interval between spikes is a consequence of the spike sorting process. Nonetheless, this is a crucial quantity regarding the difference between firing rate and clipped firing rate (see fig. 2.2), which must be taken into consideration when choosing an appropriate surrogate method.




Figure 3.2: We show the statistical measures introduced in chapter 2, i.e., C_V , C_{V2} , and FF for the recordings of both monkeys. Additionally, we present the minimal inter-spike intervals (ISI) for each neuron. The dashed lines stand for the minimal possible value of the ISI that is fixed during the spike sorting procedure (Brochier et al. (2018)). For all of these four plots, these measures are plotted against the average firing rate of the neuron unit. Each dot in the scatter plots represents a neuron. We highlight the neurons with the five highest average firing rates again in colors.

Chapter 4

SPADE

SPADE is an algorithm for the detection and evaluation of significant spatio-temporal spike patterns at a precision level of milliseconds (Torre et al. (2013); Quaglio et al. (2017); Stella et al. (2019)). It has been shown that it outperforms other methods for the detection of spatio-temporal patterns (Stella (2017); Quaglio et al. (2018)). Its development started as a tool for the detection of synchronous patterns (Torre et al. (2013)) and was extended to account for patterns with various temporal delays (Yegenoglu et al. (2016); Quaglio et al. (2017); Stella et al. (2019)). We show the workflow of SPADE in figure 4.1.




Figure 4.1: Workflow of SPADE.

The spike train data is first discretized using clipping (sec. 2.5). The clipped data is then mined for spatio-temporal patterns using the Frequent Itemset Mining (FIM) algorithm (Zaki (2004); Borgelt (2012); Picado-Muiño et al. (2013)), which yields pattern candidates that are further evaluated for significance. At the same time, using one of the dithering methods (a subcategory of surrogate methods) described and analyzed in

chapter 5, different realizations of surrogate data are created, which are equally clipped and mined yielding a corresponding p-value distribution, called the p-value spectrum (Stella et al. (2019)). Subsequently, the candidate patterns can be selected according to the p-value, corrected for multiple testing (Stella et al. (2019)). The test on the patterns detected in the original data based on those retrieved in the surrogates is called Pattern Spectrum Filtering. At the last step, the set of patterns is reduced, analyzing the mutual conditional probabilities for each pair of two patterns. This step is called pattern set reduction (Torre et al. (2013)).

In the following, we want to explain the mining technique FIM briefly and in detail, the p-value spectrum and the pattern spectrum filtering. Since the pattern set reduction does not strongly influence the evaluation of the dithering methods, we do not discuss it here; it can be found in the mentioned paper.

4.1 Definition of Spatio-Temporal Patterns

In order to detect and evaluate spatio-temporal patterns, it is necessary to have a consistent definition of it. Given a data set that contains N spike trains defined in the time range $[0, T]$, whose clipped version is written s_m^n for the n -th neuron¹, we define a pattern of neurons $\{n_1, \dots, n_z\}$ with delays $\{d_1, \dots, d_z\}$ ($d_i \in \mathbb{N}_0$; $\min(d_i) = 0$) and a number of occurrences c such that

$$\sum_{m=0}^{\lfloor \frac{T}{b} \rfloor - \max(d_i)} \prod_{i=1}^z (s_{m-d_i}^{n_i}) = c. \quad (4.1)$$

We show an example of this in figure 4.2. The bin positions of the pattern $\{m_1, \dots, m_c\}$ correspond to those of the first spikes of the pattern. We can conclude that a spatio-temporal pattern is uniquely described by the set of neurons $\{n_1, \dots, n_z\}$, the set of delays $\{d_1, \dots, d_z\}$, and the set of non-zero bins $\{m_1, \dots, m_c\}$. To simplify the classification of the patterns, we assign to each pattern a signature (z, c, d) , with z being the number of neurons, c the number of occurrences, and d the maximal delay.

4.2 Mining of spatio-temporal patterns

The Frequent Itemset Mining (FIM) is a widely used mining technique originating from market basket analysis (Zaki (2004); Borgelt (2012)). It was first applied to spike train data by Picado-Muiño et al. (2013). A very instructive overview figure can be found in Stella et al. (2019). Besides needing clipped spike trains as input and having a low cost computationally, its outstanding feature is that it identifies directly patterns which are

¹In order to facilitate, we omit here the “clip” superscript.




Figure 4.2: Spatio-temporal pattern extracted from continuous-time spike train data. Where first the step of clipping is shown and then the elements of formula (4.1). The pattern taken into consideration consists of the neurons $\{1, 2, 3, 4\}$ with the delays $\{2, 1, 1, 0\}$ and the bin positions $\{1, 3, 5\}$. We present this pattern's spikes in green. Thus its signature contains the number of neurons $z = 4$, the number of occurrences $c = 3$, and the duration $d = 2$ (the signature is thus $(4, 3, 2)$). In the first panel, we show the continuous-time representation of the spike train, where each tick represents a spike. In the second panel, the spike trains are discretized (clipped), and in each bin, we represent the presence/absence of a spike with 1/0, respectively. The clipped spike trains are shifted in the third panel to represent their alignment when considering their lags (product term in formula 4.1). Finally, the aligned spikes are summed to count the number of occurrences in the last panel.

trivially explained by others (Torre et al. (2013))². Consider that the minimal occurrence that FIM searches for patterns is set to $c_0 > 1$ and the maximal duration to $d_0 \geq 0$.

4.3 P-Value spectrum and pattern spectrum filtering

A central element of the significance analysis of patterns in SPADE is the creation of the p-value spectrum (Stella et al. (2019)). As we have seen, each pattern is characterized by a signature (z, c, d) ³, containing the number of included spikes z , the number of occurrences in the data c , and the duration, i.e., the difference in bins between the first and the last spike d . The set of pattern signatures for the data set is denoted S .

To obtain a p-value for each signature, n_{surr} surrogates are created by one of the surrogate methods described in chapter 5. An analysis with FIM gives for each surrogate j a list of pattern candidates in the data. The set of pattern signatures for this surrogate is called S_j .

From each surrogate data set a binary function is extracted:

$$f_j(z, c, d) = \begin{cases} 1 & \text{if } \exists z^* \geq z \wedge \exists c^* \geq c \text{ with } (z^*, c^*, d) \in S_j \\ 0 & \text{else} \end{cases} \quad (4.2)$$

The p-value spectrum results as an average of these binary functions over the surrogates.

$$p_v(z, c, d) = \langle f_j(z, c, d) \rangle_{\text{surr}} \quad (4.3)$$

Thus, the p-value spectrum function describes for each pattern signature (z, c, d) the probability of finding a signature with higher or equal pattern size z and with higher or equal pattern occurrence c in the data set.

An example of this three-dimensional p-value spectrum can be seen in figure 4.3⁴. With the help of this figure it can also be intuitively understood, why not only the pattern size z and its occurrence c are parameters for the p-value spectrum, but also the duration d . With increasing size z the combinations of the ways how a by-chance pattern can occur increase roughly proportional to d^{z-2} .⁵ Thus, for a pattern to be significant, it needs more occurrences for higher durations, which is accounted for in the 3-dimensional p-value spectrum (Stella et al. (2019)).

²If the set of spikes of a pattern is a subset of spikes of another pattern, then it is trivially explained by this other pattern.

³There are two versions of SPADE one where the signature contains only the size and the occurrence, and one which also contains the duration. The latter, called 3d-SPADE, was introduced in Stella et al. (2019). This paper also analyzed the advantages of the new version examining false positives and false negatives. In this thesis, we only use the 3d-version of SPADE.

⁴Readers may also refer to the corresponding figure in Stella et al. (2019).

⁵The term d^{z-2} arises from the fact that $z - 2$ spikes can be placed freely within a window of length d , since one spike has to be the first, and one has to be the last.




Figure 4.3: Three-dimensional p-value spectrum that results as a mean value from the realizations of 100 p-value spectra. Each was created by the use of 5000 Joint-ISI surrogates ($\Delta = 15$ ms) of 20 independent PPR spike trains ($\tau_r = 3$ ms) of a rate of 80 Hz and a duration of 1 s. In the three-dimensional p-value spectrum, a p-value is assigned to each signature (z, c, d) . Each subplot corresponds graphically to a pattern size z . The number of occurrences c is shown on the x -axis and the duration in bins d on the y -axis. A color then represents the p-value for each point in this three-dimensional space. On the right side, we show the corresponding logarithmic color bar. Here the reddish colors correspond to p-values from 1 to about 0.1, while in bluish tones all p-values below 10^{-4} are shown. If no value is displayed, the p-value determined by surrogates is precisely zero.

To finally evaluate the patterns for significance, the patterns of S are first reduced to those having a p-value smaller than 1: S' . The signatures of S' are then compared to a significance threshold α using a multiple testing correction typically either the Holm-Bonferroni correction (Holm (1979)) or the Benjamini-Hochberg correction (Benjamini and Hochberg (1995)).

As described above, as the last step, SPADE performs the pattern set reduction test to remove the patterns that are a by-product of the overlap of existing patterns and chance spikes. This test consists of three subtests: the subset filtering, the superset filtering, and the test for covered spikes (Torre et al. (2013)). Each of these tests can be adjusted with a parameter. The underlying question is how to compare two overlapping patterns, one with more spikes and fewer occurrences and one with fewer spikes and more occurrences.

The patterns detected by FIM and surviving the two testing steps are the final result of the SPADE method.

Chapter 5

Surrogate Methods

In order to assess the significance of spatio-temporal patterns, we have to formulate a null-hypothesis accordingly. In our context, the null hypothesis states that the patterns detected in the data are a by-chance product of the statistical properties of mutually independent spike trains (Quaglio et al. (2017); Stella et al. (2019)). To implement this null hypothesis, we employ a Monte Carlo approach that creates surrogates with similar properties as the original spike trains. Since the time-varying firing rates and other properties can only be estimated for spike trains from electrophysiological recordings, the used surrogate methods randomly displace each spike slightly from its original position¹. The approach chosen for previous studies in the context of the SPADE method is called "dithering" (Louis et al. (2010a)). The same approach was already used widely in literature and similar studies (Date et al. (1998); Nádasdy et al. (1999); Hatsopoulos et al. (2003); Shmiel et al. (2006); Stark and Abeles (2009); Louis et al. (2010b)).

We will present three dithering methods below. We will then analyze them in terms of the surrogate spike trains' statistical properties, like the resulting ISI-distribution and autocorrelation function. Based on this, we will evaluate which of the surrogate methods are most suitable for our requirements.

5.1 Uniform dithering

The method of uniform dithering was introduced in Date et al. (1998) and had become one of the standard approaches to create surrogate spike trains (Louis et al. (2010a)). In order to destroy precise temporal correlations and thus to create almost mutually independent spike trains, uniform dithering adds an increment δ to each spike time, where δ is uniformly distributed: $\delta \propto \mathcal{U}(-\Delta, \Delta)$ (Louis et al. (2010a)).

The dither parameter $\Delta > 0$, which must be chosen appropriately, determines the maximal displacement of a spike from its original position. If we write the **p.d.f.** of the

¹An overview can be found in Louis et al. (2010a).




Figure 5.1: Sketch explaining the uniform dithering method. On top, we represent the original spike train. We embody each spike by a tick (green). On the bottom, we display the corresponding dithered spike in orange. Grey dotted boxes represent the range of the interval of the uniform distribution, centered on the original spike. Note that the ranges for different spikes can intersect each other. Figure inspired by Louis et al. (2010a).

displacement of each spike, which is a normalized boxcar function as²

$$u(\delta) = \frac{1}{2\Delta} \Theta(\Delta - |\delta|), \quad (5.1)$$

the firing rate of the surrogates follows as (Louis et al. (2010b)):

$$\lambda_\delta(t) = (\lambda \circledast u)(t). \quad (5.2)$$

In general, this will smooth the underlying firing rate of a spike train (Louis et al. (2010b)).

The cross-correlation ψ_{cc} between a spike train $s(t)$ and its surrogate spike train $s_\delta(t') = s(t' + \delta)$, is described by

$$\psi_{cc}(t, t') := \langle \langle s(t) s_\delta(t') \rangle_s \rangle_\delta \stackrel{(2.7)}{=} \langle \psi(t, t' + \delta) \rangle_\delta = \int d\delta u(\delta) \psi(t, t' + \delta). \quad (5.3)$$

The quantity is important to measure in how far the surrogate spike trains differ from the original ones. Thus, we introduce as a measure of similarity the rate of spikes which are in the same bin for the original and the surrogate spike trains: $\lambda_b(t)$. This is evaluated by integrating over the cross-correlation function for one bin size $b > 0$ as

$$\begin{aligned} \lambda_b(t) &:= \frac{1}{b} \left\langle \left\langle \int_{-b/2}^{b/2} d\tau \int_{-b/2}^{b/2} d\tau' s(t + \tau) s(t + \tau' + \delta) \right\rangle_s \right\rangle_\delta \\ &\stackrel{(5.3)}{=} \frac{1}{b} \int_{-b/2}^{b/2} d\tau \int_{-b/2}^{b/2} d\tau' \psi_{cc}(t + \tau, t + \tau') \end{aligned} \quad (5.4)$$

²The corresponding average operator is $\langle \dots \rangle_\delta = \int d\delta u(\delta) \dots$.

For renewal processes the cross-correlation follows directly by plugging equation (2.38) into equation (5.3):

$$\psi_{cc}(t, t') = \lambda u(t - t') + \lambda(\lambda_c \otimes u)(t - t'). \quad (5.5)$$

Thus, the cross-correlation between original data and its surrogate contains two-terms. The first one being a boxcar function, and the second one being the convolution of the conditional firing rate λ_c (see eq. (2.34)) and the p.d.f. of the uniform dithering.

From now on, for the sake of simplicity, we will concentrate on renewal processes only. As we observed already in section 2.4.1, renewal processes exhibit a constant rate. For this reason, by plugging equation (5.5) into (5.4), we see that the rate of spikes in the same bin is thus constant over time containing two terms:

$$\begin{aligned} \lambda_b &= \frac{1}{b} \int_{-b}^b d\tau (b - |\tau|) \left(\lambda u(\tau) + \lambda(\lambda_c \otimes u)(\tau) \right) \\ &= \frac{\lambda b}{2\Delta} + \frac{\lambda}{b} \int_{-b}^b d\tau (b - |\tau|) \left((\lambda_c \otimes u)(0) + \mathcal{O}(\tau^2) \right) \\ &= \lambda b \left(\frac{1}{2\Delta} + (\lambda_c \otimes u)(0) \right) + \mathcal{O}(b^3). \end{aligned} \quad (5.6)$$

Note that the linear term of the expansion of $(\lambda_c \otimes u)(\tau)$ vanishes due to the symmetry of both λ_c and u .

The first term is proportional to the probability that a spike will be displaced within the same bin in which it previously was: this probability is $\frac{b}{2\Delta}$. A trade-off arises when choosing the dither parameter Δ appropriately. On the one hand, the firing rate should not be too strongly flattened (see eq. (5.2)) and as well the firing rate of the clipped surrogate spike trains (see eq. (5.14)), discussed later on, should be as high as possible, while on the other hand λ_b should be small. If the λ_b is high, not enough spikes may be displaced to destroy temporal correlations. As a result of this reasoning, if the bin size b is varied, Δ should be changed accordingly such that the ratio $\frac{b}{\Delta}$ stays constant.

The second term stands for spikes displaced into the bins where other spikes had been in the original spike train since it originates from the second part of the autocorrelation (eq. (2.38)). This term impacts more for high firing rates because of the conditional firing rate λ_c scaling with the firing rate λ .³

In the appendix B.2, we derive closed forms for the expression of the rate of spikes falling in the same bin λ_b for the case of the PPR and the gamma process (where we make use of the effective firing rate $\tilde{\lambda}$ ⁴ and the incomplete gamma function $\bar{\Gamma}$ ⁵):

³The asymptotic behavior of $\lambda_c(\tau)$ is that it tends towards the firing rate λ for $\tau \rightarrow \pm\infty$ because the presence of a spike does not affect the probability of having a spike infinitely far away from it. Comparing to figure 5.2, we conclude that it need only a few steps of the mean ISI, i.e., $1/\lambda$, that $\lambda_c(\tau)$ is very close to λ .

⁴Effective firing rate: $\tilde{\lambda} = \frac{\lambda}{1 - \lambda\tau_r}$ (see eq. (2.60))

⁵Incomplete gamma function: $\bar{\Gamma}(\gamma, x) := \int_x^\infty ds s^{\gamma-1} e^{-s}$ (see eq. (2.68))

$$\lambda_b^{\text{PPR}} = \frac{\lambda b}{2\Delta} \left(1 + 2 \sum_{n=1}^{\lfloor \frac{\Delta}{\tau_r} \rfloor} \frac{\bar{\Gamma}(n, \tilde{\lambda}(\Delta - n\tau_r))}{\Gamma(n)} \right) + \mathcal{O}(b^3), \quad (5.7)$$

$$\lambda_b^\gamma = \frac{\lambda b}{2\Delta} \left(1 + 2 \sum_{n=1}^{\infty} \frac{\bar{\Gamma}(n\gamma, \gamma\lambda\Delta)}{\Gamma(n\gamma)} \right) + \mathcal{O}(b^3). \quad (5.8)$$

For the Poisson process, we can solve λ_b without the approximation of $(\lambda_c \otimes u)(\tau)$, as applied in eq. (5.6), since $\lambda_c = \lambda$, an effect of the memoryless property (see eq. (2.57)). This equality plugged into eq. (5.6) yields

$$\lambda_b^{\text{Pois}} = \lambda b \left(\frac{1}{2\Delta} + \lambda \right). \quad (5.9)$$

The autocorrelation function of the surrogate spike trains is closely related to the ISI distribution of the surrogates. As stated at the beginning of this section, surrogate spike trains should be almost mutually independent and should approximately conserve the original spike trains' statistical features. We would see this conservation if the surrogates' autocorrelation function is close to that of the original point process.

In the appendix B.1, we show how this autocorrelation for uniform dithered surrogates can be derived, which is written as

$$\begin{aligned} \psi_\delta(t, t') &= \lambda_\delta(t) \delta(t - t') + \int ds \int ds' \phi(s, s') u(t - s) u(t' - s') \\ &\text{with } \phi(s, s') := \psi(s, s') - \lambda(s) \delta(s - s'). \end{aligned} \quad (5.10)$$

It contains the smoothed firing rate (see eq. (5.2)) and a term that follows from two convolutions with the boxcar function of the uniform dithering.

For the case of renewal processes, we can use the autocorrelation already formulated in chapter 2, and obtain the surrogate autocorrelation function as:

$$\psi_\delta(t, t') = \lambda \delta(t - t') + \lambda(\lambda_c \otimes u \otimes u)(t - t'). \quad (5.11)$$

We show different autocorrelation functions for PPR and gamma processes and their uniformly dithered surrogates in figure 5.2. Additionally, we vary the refractory period τ_r for the PPR and the shape parameter γ for the gamma process. The initial processes are in full lines, while dotted lines represent the surrogates. Most importantly, we see in the figure that the uniform dithering does not conserve the spike train's refractory period. Besides the refractory period, the autocorrelation function is, in general, not conserved after uniform dithering in the case of these processes. Considering that the solid blue line in the subplot on the right corresponds to the autocorrelation of the




Figure 5.2: Autocorrelation function before (solid line) and after uniform dithering (dashed line), not displaying the central delta-function. On the left, we show PPR processes with different refractory periods (in different colors). On the right, we see gamma processes with different shape factors presented in different colors. Note that the gamma process with $\gamma = 1$ is the Poisson process. Since $\psi(t, t')$ depends only on the difference between t and t' , we show on the x -axis this difference in ms. On the y -axis, we show the autocorrelation function over the firing rate. We take this ratio, such that the baseline corresponds in both cases to the firing rate. The firing rate of 80 Hz is at the upper end of firing rates of neuronal recordings (see chap. 3). We choose such a high firing rate since we observe high numbers of false positives for uniform dithering for high rates (see chap. 6). The dither parameter of 15 ms used here is a typically chosen value (Quaglio (2019)).

Poisson process, it is possible to state that surrogates created by uniform dithering tend to have *more Poisson-like properties* than the original spike trains since the dotted lines tend to get closer to the blue line.

In order to measure the firing rate of the clipped spike train after applying uniform dithering, it is necessary to approximate the ISI distribution for small τ . We will, therefore, simplify by approximating the resulting surrogate spike trains to be renewal. To approximate like this is only possible for the regime of very small τ , especially regarding the negative serial correlations that arise from uniform dithering (see fig. 5.11). In this way, we will be able to calculate the ISI distribution from the autocorrelation function (which is directly related to the conditional firing rate) of the surrogate spike trains.

In the appendix B.3, we show how to obtain the ISI distribution of a renewal process




Figure 5.3: Clipped firing rate for PPR and gamma processes before and after uniform dithering (solid/dashed line). We use a dither parameter of $\Delta = 15$ ms, and a bin size of $b = 3$ ms. Note that the gamma process with $\gamma = 1$ is the Poisson process. The different colors correspond to different configurations of the PPR spike trains (left) and the gamma spike trains (right). While the x -axis shows the initial firing rate in Hz, the y -axis shows the ratio of the clipped firing rate over the initial firing rate.

after uniform dithering in the small τ limit. The result is

$$p_\delta(\tau) = (\lambda_u - (\lambda_u)^2\tau + \mathcal{O}(\tau^2)) \Theta(\tau) \text{ with } \lambda_u := (\lambda_c \circledast u \circledast u)(0). \quad (5.12)$$

The second convolution of the conditional firing rate at $\tau = 0$, which we call here λ_u can be understood as

$$\lambda_u = \lim_{\epsilon \rightarrow 0} \frac{1}{\epsilon} \text{Prob}(\text{spike in } [t, t + \epsilon] \text{ after UD} \mid \text{other spike at } t \text{ after UD}) \forall t > 0. \quad (5.13)$$

Thus, it reflects the probability in the surrogates that a spike follows directly after another one.

The clipped firing rate follows (by using eqs. (2.37) and (2.74))

$$\lambda_{\text{clip}} = \lambda \left(1 - \frac{1}{2}\lambda_u b + \frac{1}{6}(\lambda_u b)^2 \right) + \mathcal{O}(b^3). \quad (5.14)$$

Values for λ_u for PPR and gamma processes can be found in appendix B.2. We show the dependence of the clipped firing rate on the firing rate in figure 5.3. Both for the

PPR and the gamma spike trains (with $\gamma > 1$), we observe a substantial decrease of the clipped firing rate, i.e., the solid lines in figure 5.3 representing the original spike trains lay much higher than the dashed lines representing the surrogates. Thus, using uniform dithering as the surrogate method yields a *reduction in the spike count*, and as we will see in chapter 6, this is a *source of false positives* that need to be taken into account carefully.

5.2 Uniform dithering with refractoriness




Figure 5.4: Scheme showing how to achieve the conservation of the absolute refractory period when transitioning from uniform dithering to uniform dithering with refractoriness. The left part is the same as in figure 5.1. In the upper blue line, every green tick stands for a spike of an original spike train, while the orange ticks of the lower line represent the spikes of the surrogate. The boxes refer to the space in which the spike can be possibly displaced via the uniform dithering. For the UDR, shown on the right, these boxes are changed that they do not overlap anymore. Not only do they not overlap, but they are adjusted such that the interval between two spikes in the surrogate cannot be smaller than the refractory period τ_r . Figure inspired by Louis et al. (2010a).

As we have seen in the previous chapter, the firing rate decreases strongly due to the interference between clipping and uniform dithering. To prevent this, we introduce an adaptation of the uniform dithering that we call **uniform dithering with refractoriness** (UDR).⁶

The basic idea is that we adjust the dithering range for each spike in such a way that *the absolute refractory period is conserved* (see fig. 5.4). We have seen in figure 2.2 how the clipped firing rate depends on the refractory period. Essentially, the closer the refractory period to the bin size, the closer the clipped firing rate to the firing rate. Note that by the way, spike trains are recorded extracellularly and spike sorted, an absolute refractory period always exists (see sec. 3.1). Therefore, preserving the refractory period will, for all spike trains, reduce the difference between clipped and initial firing rate.

⁶The reader can find the implementation of this algorithm as part of the software package elephant.

For this purpose, we estimate the refractory period of the spike train data as the minimum of inter-spike intervals and an initial (biological) guess $\tau_r^{\text{init}} > 0$.

$$\tau_r^{\text{est}} = \min(\min(t_i - t_{i-1}), \tau_r^{\text{init}}) \quad (5.15)$$

For the sake of conserving the refractory period between spike i and spike $i + 1$, the maximal displacement of spike i , i.e., Δ_+^i , has to be such that $t_i + \tau_r^{\text{est}} + \Delta_+^i \leq t_{i+1}$. Since the UDR should go in the limit of low firing rates back to the uniform dithering, Δ_+^i can be maximally Δ . Subsequently, we set Δ_+^i to:

$$\Delta_+^i := \min((t_{i+1} - t_i - \tau_r^{\text{est}}), \Delta). \quad (5.16)$$

In the same way, we define Δ_-^i as:

$$\Delta_-^i := \min((t_i - t_{i-1} - \tau_r^{\text{est}}), \Delta). \quad (5.17)$$

The increment δ added to spike i is distributed as $\delta \propto \mathcal{U}(-\Delta_-^i, \Delta_+^i)$:

$$u_r^i(\delta) := \frac{1}{\Delta_+^i + \Delta_-^i} \Theta(\Delta_+^i - \delta) \Theta(\delta + \Delta_-^i). \quad (5.18)$$

It is important to note that every spike is updated individually, where the order of the spikes to be updated is a random sequence.

For low firing rates, the difference of the adapted version to the original uniform dithering is small. In contrast, for high firing rates, we can expect the spikes to be displaced less, on average, since every spike is tied between the spikes surrounding it.

The analytical evaluation of the quantities we introduced for uniform dithering is not possible to perform here since every spike is displaced, depending on the positions of its neighbors. For this reason, we would need to introduce more advanced point process models allowing the incorporation of serial ISI correlations to study the behavior of the UDR analytically.

Nonetheless, it is possible to say that non-negligible serial correlations arise when creating surrogate spike trains from renewal processes. Instead of calculating analytically, we refer to the later section 5.4, where the autocorrelation and cross-correlation functions and their deduced quantities are compared numerically across the surrogate methods.

5.3 Joint-ISI dithering

The joint-ISI dithering is a method that was introduced by George Gerstein in 2004 (Gerstein (2004)) and further developed in Louis et al. (2010b). Gerstein introduced it as a method that (approximately) conserves firing rate profile and ISI distribution.




Figure 5.5: The figure shows the joint-ISI distribution of a spike train and the displacement of a spike within the joint-ISI space. The joint-ISI distribution displayed here is the analytical joint-ISI distribution for a gamma process of a firing rate of 50 Hz. Since it is a renewal process, the joint-ISI distribution corresponds to the product of two ISI distributions (see eq. (2.32)). In both x- and y-axes, we show ISIs in ms, i.e., the interval to the previous τ and the interval to the next τ' . The color code displays the values of the joint-ISI distribution. A spike is shown with intervals of 25 ms to the previous spike and 20 ms to the next spike. This spike can be displaced on the dotted line, limited by $\pm\Delta$. On this line, the spike is displaced according to a one-dimensional projection of the two-dimensional joint-ISI distribution. Figure inspired by Louis et al. (2010b).

An essential difference to the previous methods is that the spikes are not displaced according to a uniform distribution. Instead, a spike is displaced based on the intervals to the previous and subsequent spike. The change in these intervals happens according to the joint-ISI distribution.

For simplicity, we will first introduce the displacement procedure for a continuous joint-ISI distribution.

Given a spike train with joint-ISI distribution $f(\tau, \tau')$, the displacement of the spike i at time t_i will depend on the time points t_{i-1} and t_{i+1} , and it is defined as⁷

$$u^i(\delta) \propto f((t_i + \delta) - t_{i-1}, t_{i+1} - (t_i + \delta))\Theta(\Delta - \delta). \quad (5.19)$$

How the displacement of spikes looks like in the joint-ISI space can be seen in figure 5.5. It mostly conserves the ISI-distribution (see sec. 5.4) and the absolute refractory

⁷The original version of the joint-ISI dithering applied a square root to the joint-ISI distribution (Gerstein (2004); Louis et al. (2010b)). In the appendix sec. B.4, we argue why this is only reasonable for small Δ . Inside the implementation that we added to the software package elephant, users still can use the original version.

period by moving each spike by the corresponding section of the joint-ISI distribution. Conserving the refractory period is a result of the constraint:

$$f(\tau, \tau') = 0 \text{ for } \tau < \tau_r \text{ or } \tau' < \tau_r, \quad (5.20)$$

which together with eq. (5.19) does not allow moving a spike inside a refractory period of another spike. Additionally, we first update the even- and then the odd-numbered spikes, which is computationally advantageous⁸.

As we discussed in sec. 2.3 for finite time spike trains, it is necessary to use a binned version of the Joint-ISI distribution applying adapted Gaussian filtering. In this binned fashion, the probability of displacing i -th spike by d bins of size b becomes:

$$u_d^i \propto \Theta(\lfloor \frac{\Delta}{b} \rfloor - d) f_{\lfloor \frac{t_i - t_{i-1}}{b} \rfloor + d, \lfloor \frac{t_{i+1} - t_i}{b} \rfloor - d}^\sigma. \quad (5.21)$$

Additionally, we truncate the joint-ISI histogram into a triangle with side length W , i.e., $f_{n,n'}^\sigma$ is only defined for $m + m' \leq \lfloor \frac{W}{b} \rfloor$ (lower triangle of the joint ISI distribution) since the estimation of the joint-ISI distribution tends to be worse the larger the inter-spike intervals. Due to the triangle structure, more bins would need to be filled for higher ISIs. In our implemented approach, if a set of two inter-spike intervals is outside this triangle, the corresponding spike in the middle is dithered according to the UDR protocol (see sec. 5.2).

To create a suitable null hypothesis using joint-ISI surrogates, we need to choose suitable parameters for the dither parameter Δ , the kernel width σ , the truncation limit W , and an initial guess for the refractory period τ_r^{init} ⁹. Further, we select the bin size b with which we create the joint-ISI histogram. This bin size is not necessarily the same as the one used in SPADE, usually smaller. How to choose all these parameters depends on the statistics of the spike train data.

We will present the performance of the joint-ISI dithering regarding auto- and cross-correlation functions in the next section.

5.4 Comparison of surrogate methods

Comparing the different surrogate methods is based on examining different measurable quantities before and after applying the surrogate methods. For this purpose, we test the three above introduced methods on PPR and gamma processes. Onto these, we then apply uniform dithering and UDR, both with a dither parameter of $\Delta = 15$ ms. Finally, we use joint-ISI dithering with a dither parameter of $\Delta = 15$ ms, a kernel width of $\sigma = 1$ ms, an initial refractory period of $\tau_r^{\text{init}} = 4$ ms, and a bin size of $b = 1$ ms.

⁸This is in contrast to the procedure of Louis et al. (2010b).

⁹The initial refractory period τ_r^{init} is used in eq. (2.28).

5.4.1 Firing rate profile

First of all, we analyze how the dithering methods affect the firing rate profile. To this end, we generate data, the firing rate profile representing a step function that first has a firing rate of 10 Hz for 50 ms and then a firing rate of 80 Hz for another 50 ms. These are the same parameters used for a similar analysis in Louis et al. (2010a). We show the resulting firing rate profile for each surrogate method in figure 5.6. Note that the firing rate change does not depend on the further statistical properties of the spike train. For the uniform dithering, we can see that the theoretically determined change in the firing rate is confirmed (see eq. 5.2). The deviations from the result of the uniform dithering for UDR and joint-ISI dithering are very small. Interestingly, the firing rate profile that is yielded by UDR is significantly lower for the interval between 60 ms and 70 ms, which can be explained as a consequence of the asymmetric displacement protocol (see eq. 5.18). Nonetheless, we conclude that all three methods *smooth the firing rate similarly*.

5.4.2 Cross-correlation

The second quantity to analyze is the cross-correlation function, expressing the similarity between the original artificial data and its surrogates. We show it in figure 5.7. For uniform dithering, we see that for PPR and gamma spike trains, the cross-correlation has its maximum at $\pm\Delta = \pm 15$ ms. At the same time, when using UDR or joint-ISI dithering on the PPR spike trains, it is possible to see a reflection of the refractory period in the cross-correlation. There are peaks in the cross-correlation at ± 3 ms¹⁰ and a lower peak at 0 ms. These peaks are a consequence of spikes whose dithering range is limited by the refractory period. For gamma processes, similar behavior between UDR and joint-ISI dithering can be observed, with the difference that the UDR has a peak for 0 ms, while for joint-ISI dithering at 0 ms it has a local minimum.

Regarding the rate of spikes in the same bin for the original data and the surrogate, which we have introduced as λ_b , we can deduce from the cross-correlation that for PPR spike trains UDR and joint-ISI dithering will perform similar and better as the uniform dithering. While for the gamma process, the UDR will perform worse than the other two, a result of small ranges of the increments (see eq. (5.18)) if the average ISI becomes smaller than the dither parameter Δ .

Quantitatively λ_b is displayed in figure 5.8. We see there that the ratio λ_b/λ increases monotonously for all surrogate methods and is higher than the corresponding quantity calculated for a spike train generated independently, with the same characteristics. Further, we observe that for the PPR spike trains (left figure), UDR and joint-ISI dithering have almost identical values, and they are lower than the value for the uniform dithering. Instead, gamma spike trains attain the highest value in the case of UDR and lowest for the joint-ISI dithering. Since a high value corresponds to effectively a low number of spikes displaced, we conclude that, regarding λ_b , the *joint-ISI dithering* performs best.

¹⁰It is not at the refractory period of 2.5 ms since the cross-correlation function uses a bin width of 1 ms.




Figure 5.6: The figure shows the firing rate profile before and after dithering. On the left side, this is shown for PPR spike trains ($\tau_r = 2.5$ ms) and on the right side for gamma spike trains ($\gamma = 2$). For both, we generate 10^5 realizations, and we dither each realization once with each surrogate method. To omit boundary effects, we create them for a domain of -15 ms to 115 ms, but we only show the firing rate between 0 ms and 100 ms. To obtain the firing rate, we count the number of spikes per bin (bin width of 1 ms). The resulting firing rate in Hz is shown as a function of time in ms. We see it for the artificial data itself (art. in blue) and their corresponding results of the uniform dithering (UD in orange), the uniform dithering with refractoriness (UDR in green), and the joint-ISI dithering (JISI in red). The used parameters for the surrogate methods, we describe in sec. 5.4.




Figure 5.7: The cross-correlation function of artificial data against its surrogate is displayed here divided by the firing rate. Thus, the baseline is the firing rate, i.e., 80 Hz. On the left panel, it is shown for PPR spike trains ($\tau_r = 2.5$ ms) and on the right for gamma spike trains ($\gamma = 2$). Due to the time translation invariance for renewal processes, we show a projection of $\psi_{cc}(t, t')$ depending only on the difference $t - t'$ given in ms. The results for the three surrogate methods are differently colored, i.e., uniform dithering (UD), uniform dithering with refractoriness (UDR), and joint-ISI dithering (JISI). For each spike train type, we generate a 6250 s long spike train that we dither using each of the three methods. The cross-correlation function is obtained with a bin size of 1 ms.




Figure 5.8: The plot shows the ratio of the rate of spikes in the same bins for the original spike train and surrogate over the initial firing rate shown as a function of the initial firing rate in Hz. On the left panel, it is for PPR spike trains, while on the right panel for gamma processes. The solid blue lines show the value of this ratio if instead of a surrogate, we create another independent spike train of this type. The other dashed lines show how many spikes are in the same bin, if the spike train is dithered with uniform dithering (UD), uniform dithering with refractoriness (UDR) or joint-ISI dithering (JISI). We obtain these values on spike trains with an expected spike count of $5 \cdot 10^5$ spikes, the lowest firing rate being 5 Hz, which is increased in steps of 5 Hz. It is noteworthy that for the discretization here, the spike trains were only binned and not clipped (see sec. 2.5).




Figure 5.9: The ISI distributions for PPR processes ($\tau_r = 2.5$ ms, left) and gamma processes ($\gamma = 2$, right) are shown and of their surrogates. The x -axis corresponds to the ISI τ in ms and the y -axis to the ISI-distribution $p(\tau)$ in Hz. The values for the artificial data (art.) are in blue, the ones for the surrogates methods in orange (uniform dithering, UD), green (uniform dithering with refractoriness, UDR), and red (joint-ISI dithering, JISI). The generated data is the same as for fig. 5.7.

5.4.3 ISI-distribution

Since its ISI distribution entirely describes a renewal process, we examine the resulting ISI distribution after the application of all dithering methods in figure 5.9. The joint-ISI dithering should by construction conserve the ISI-distribution. We observe that this is mostly fulfilled, but we also see a small reduction of the ISI distribution for small τ . This reduction for small τ is compensated by a shift of probability mass into the range from 10 to 20 ms. Thus, the coefficient of variation will also increase slightly. For the uniform dithering with refractoriness, we observe that applied on the PPR spike train, it almost exactly preserves the ISI distribution, while for the gamma spike train, the yielded ISI distribution has low similarity to the original distribution. For uniform dithering in both cases, the ISI distributions appear to be similar to exponential decay, and the original process cannot be distinguished anymore. Thus, we conclude that the *joint-ISI dithering* is the method of choice given the constraint that the ISI distribution is part of the statistical properties mentioned in the null hypothesis.

5.4.4 Autocorrelation function

Closely linked to the ISI-distribution is the autocorrelation function (see eqs. (2.33) & (2.38)). While the ISI-distribution shows when the next spike happens, the autocorrelation function shows when other spikes occur. One of the information that it contains is the interval after a spike on which it effectively has no influence anymore on the probability of the next spike, i.e., when the autocorrelation converges to the baseline. Also, it shows when it becomes more likely to have a spike after a certain interval. We can observe this in figure 5.10. The autocorrelation for the surrogate spike trains of the joint-ISI dithering applied on PPR and gamma spike trains is higher than the baseline in a range of approximately 10 to 20 ms. The same holds for the UDR applied on PPR spike trains and is a sign of arising positive serial correlations (van Vreeswijk (2010)). We see that both methods can come closer to the autocorrelation of the artificial data as the uniform dithering. It is difficult to conclude from this since neither the arising of serial correlations nor the deviation from the autocorrelation function is a desired property of the surrogate method.

Regarding the serial correlations, whose coefficients are defined as (Perkel et al. (1967))

$$\xi_k = \frac{\langle \tau_{i+k} \tau_i \rangle - \langle \tau_i \rangle^2}{\langle \tau_i^2 \rangle - \langle \tau_i \rangle^2} \text{ with } \tau_i = t_{i+1} - t_i, \quad (5.22)$$

we can measure the dependence of inter-spike intervals on their predecessors. In a renewal process, there are no serial correlations since every ISI is drawn independently from the same distribution $p(\tau)$ (see sec. 2.4.1). However, as we can see in figure 5.11, serial correlations arise using the dithering methods. For uniform dithering, these are strong negative serial correlations, i.e., after having a short ISI, it is more likely to get a long ISI and vice versa. For UDR, instead, this depends on the spike train model, whether it is a PPR or a gamma process. The joint-ISI dithering shows in both cases positive serial correlations.

It is still debatable if surrogate spike trains should have rather positive or negative serial correlations. Indeed, they should be as small as possible. In this regard, the *UDR* and the *joint-ISI dithering* have more desirable properties than the uniform dithering.

5.4.5 Clipped firing rate

The last analyzed quantity is the clipped firing rate. To create a null-hypothesis suitable for conservative significance tests, the clipped firing rate of the surrogates should be close to that of the data. The comparison across the three dithering methods can be found in fig. 5.12. First of all, we can state that in the case of uniform dithering, the clipped firing rate is much lower than the initial firing rate of the data. This reduction is an effect of the uniform dithering destroying the absolute and the relative refractory period, which corresponds to the autocorrelation being unequal to zero at a zero time lag (see 5.10). Thus, the use of uniform dithering could yield having a null-hypothesis with a




Figure 5.10: We show the autocorrelation function for PPR spike trains (left panel, $\tau_r = 2.5$ ms) and gamma spike trains (right panel, $\gamma = 2$) and its corresponding surrogates, omitting the central delta-peaks. Also, using its time invariance, it is shown only the one-dimensional projection depending on the difference $t - t'$ in ms. To provide a better understanding of the baseline, we divided the autocorrelation function by the firing rate of 80 Hz. Thus, the baseline corresponds to the firing rate. We chose different colors for the artificial data (art., blue) and the surrogate methods, which are uniform dithering (UD, orange), uniform dithering with refractoriness (UDR, green), and joint-ISI dithering (JISI, red). The generated data is the same as for fig. 5.7.




Figure 5.11: The figure shows the first five orders of serial correlation coefficients for PPR spike trains ($\tau_r = 2.5$ ms, left) and gamma spike trains ($\gamma = 2$, right) after usage of a dithering method. The dithering methods are uniform dithering (UD, orange), uniform dithering with refractoriness (UDR, green), and joint-ISI dithering (JISI, red). We show the order of the coefficients on the x -axis and the coefficients on the y -axis. The data corresponds to artificially generated spike trains of a rate of 80 Hz in a time window of 10^5 s.




Figure 5.12: In this figure, we show the ratio of clipped firing rate over the firing rate for artificial data and its surrogates as a function of the firing rate. The clipping here, we perform with a bin size of 3 ms. We show the difference between PPR spike trains (left; $\tau_r = 2.5$ ms) and gamma spike trains (right; $\gamma = 2$). Also, the discrepancy can be seen between the results for the artificial data (art.) and the surrogate methods: uniform dithering (UD), uniform dithering with refractoriness (UDR), and joint-ISI dithering (JISI). The firing rate in Hz is on the x -axis and the ratio $\lambda_{\text{clip}}/\lambda$ on the y -axis. The generated data used are the same as for fig. 5.8.

significantly lower spike count. Instead, UDR fits very well with the firing rate of the artificial data for PPR spike trains, regarding gamma processes, the clipped firing rate is slightly higher than for the uniform dithering. For the gamma processes, we can explain this by the fact that there is no absolute refractory period to preserve. The high clipped firing rate for the gamma spike trains themselves is an effect of the tail of low ISI in the ISI distribution. This ISI-distribution is not preserved using the UDR (see fig. 5.9), and thus the clipped firing rate related to the UDR does not differ strongly from that of the uniform dithering. The *joint-ISI dithering* is in both cases of artificial spike trains close to the clipped firing rate of the data itself and thus seems to be in this regard an appropriate choice.

In the chapter on point processes, we have already discussed that the firing rate is proportional to the expected value of the spike count (see sec. 2.2.3). We determine the variance of the spike count's distribution using the autocovariance (see sec. 2.2.3). Also, for the clipped spike count, the clipped firing rate is proportional to the expected value, but the spike count of a single realization is not always equal to the expected value.

If we create surrogates for spike trains from electrophysiological data, only their

spike count is known, while the firing rate can only be estimated. Therefore, it is worth discussing the clipped spike count, in addition to the clipped firing rate. There is both a clipped spike count, which results solely from the clipping of the spike train, and the distribution of the clipped spike counts for the clipping of the surrogate spike trains since we generate several surrogates for each spike train, whose clipped spike counts can vary.

In figure 5.13, we now show the clipped spike counts of the electrophysiological spike trains of the reach-to-grasp data, and the mean value of the clipped spike counts for the surrogates. At first, we observe that the clipped spike count tends to decrease for higher firing rates. However, no clear function can be identified, which was to be expected, given the differences in the minimum ISIs (see fig. 3.2). Furthermore, the clipped spike count of uniform dithering is considerably lower than that of the other surrogate methods. However, no apparent differences are observable between the UDR and the joint-ISI dithering, sometimes one value is higher than the other. One reason for this is that the reach-to-grasp data does not exactly match either of the two spike train models analyzed. For most neurons, the clipped spike counts of these two surrogate methods are close to those of the spike train data. Nevertheless, we notice a considerable difference for some neurons, such as the neurons with the highest firing rates for both monkeys. Thus, we conclude that in terms of clipped spike count, *UDR* and *joint-ISI* are much closer to the original data than the uniform dithering, but they do not match completely.

5.4.6 Conclusion

We have investigated different statistical properties of surrogate spike trains. First of all, we noticed that the change in the firing rate profile is very similar for all three dithering methods. Although using the uniform dithering with refractoriness, we notice that it is more likely to displace spikes only a little with gamma processes. Concerning the resulting autocorrelation of the spike trains, there are significant differences. Only joint-ISI dithering manages to keep the ISI distribution approximately the same for both analyzed spike train types. With UDR, this is also true for PPR spike trains. Regarding the autocorrelation function, the joint-ISI dithering and the UDR are closer to the underlying function than the uniform dithering. It is noteworthy that serial correlations are measurable for all methods.

At last, and probably most importantly, we have analyzed the clipped firing rate, i.e., the effective firing rate after the discretization of time. There is a clipped firing rate both for a spike train type generically and for its surrogates. The difference between these clipped firing rates is particularly significant for uniform dithering. This reduction of the firing rate can lead to fewer spikes in the surrogates than in the original spike train data and should, therefore, be avoided. We could also observe a similar but smaller firing rate reduction in the combination of uniform dithering with refractoriness and gamma spike trains. When analyzing the reach-to-grasp data for the clipped spike count, the UDR and the joint-ISI dithering perform similarly and significantly better than the uniform dithering. Overall, concerning the clipping aspect, we recommend *joint-ISI dithering*, where the surrogates have a similar clipped spike count as the underlying spike trains.

The only disadvantage besides needing more computational time is that the estimate




Figure 5.13: Ratio of clipped spike count over the spike count for two data constellations of the reach-to-grasp data and its surrogates as a function of the average firing rate, i.e., the spike count over the recording duration. We show the spike count ratio for each neuron with a dot for clipping with a bin size of 3 ms. In the left and the right panel, we show data for two different monkeys. Also, the difference can be seen between the results for the original data (orig., blue) and the surrogate methods: uniform dithering (UD, orange), uniform dithering with refractoriness (UDR, green), and joint-ISI dithering (JISI, red). For the surrogates, it corresponds to a mean value of 200 surrogate realizations. The average firing rate in Hz is on the x -axis, and the ratio of clipped spike count over the spike count on the y -axis.

of the joint-ISI distribution is inadequate if the spike train has a low spike count. Hence, it is to apply only for a reasonable number of recorded trials or if the spike train duration is sufficiently long.

The following chapter analyzes test data with SPADE using these three different surrogate methods to evaluate false positives and false negatives and to link this evaluation to the quantities described above.

Chapter 6

False Positives and False Negatives

The analysis of false positives and false negatives is of central importance for evaluating a statistical method. In the case of SPADE, false positives and false negatives have been studied already (Quaglio et al. (2017); Stella et al. (2019)) by generating and analyzing patterns in Poisson spike trains. Here we want to extend this evaluation to PPR and gamma spike trains, focusing in particular on the comparison across the different surrogate techniques presented earlier in this thesis.

While the analysis of independently generated data can examine false positives for SPADE, an evaluation in terms of false negatives is more challenging to achieve. For this reason, we will first turn to the analysis of false positives. We restrict ourselves to stationary spike trains, whose firing rates do not vary across neurons. In a second step, we will introduce how to inject spatio-temporal patterns into independently generated data to see how the rate of false negatives depends on the surrogate method.

6.1 False positives

As we have discussed in chapter 4, the null hypothesis against which to test is that the detected spatio-temporal patterns are a by-chance product of mutually independent spike trains. The procedure is such that mined patterns are pooled by their signature (number of spikes, number of occurrences, duration), and each group of pattern candidates with the same signature can individually reject the null hypothesis. The confidence level is adjusted utilizing multiple testing. All patterns rejecting the null hypothesis are classified as statistically significant.

6.1.1 Ground truth null hypothesis

If SPADE analyzes a data set of independently generated spike trains, we should obtain a false positive rate corresponding to the confidence level. Possible reasons for false positives can be: A) the deviation of the p-value spectrum estimate to the underlying

one, if the surrogate method used is inadequate for the type of data at hand, and B) how technically the number of tests is calculated when correcting for multiple testing.

According to the previous chapters, we will deal with PPR ($\tau_r = 2.5$ ms) and gamma ($\gamma = 2$) spike trains. Thereby, we generate 100 data sets of 20 spike trains, each with firing rates ranging from 10 to 100 Hz (with steps of 10 Hz). Due to a large amount of computing time required (Stella et al. (2019)), we have limited the duration of the spike trains to one second. We used these parameters for SPADE: a bin size of 3 ms, a confidence level of 5%, and we set the maximum pattern duration d_0 to 5 bins (i.e., 15 ms). Furthermore, we looked only at patterns containing 2 to 5 spikes, allowing only one spike per neuron¹. In order to create the ground truth p-value spectrum, we generated the set of 20 spike trains again 5000 times. These data sets were then used as surrogate datasets to calculate the p-value spectrum. Additionally, p-values are corrected for multiple testing with the Holm-Bonferroni approach.

One could argue now that the spike counts of the ground truth surrogates do not match those of the generated data precisely, as is the case when using the dithering methods. This arguing corresponds to whether the same statistical properties required by the null hypothesis (see chap. 5) mean the same firing rate or the same spike count. While the first option is closer to the notion of statistical property, the second option is usually preferred because it is easier to measure. Even if we choose the surrogates' spike count to be similar to that of the original spike trains, we can still use these ground truth surrogates because the spike count distribution is the same, and the deviations can be neglected if we use enough spike trains, in our case 20. However, this is only possible because the individual spike trains' identity has no meaning due to the pooling of the pattern signatures.

In figure 6.1, the false positive rate is shown using this ground truth null hypothesis. We consider a false positive any pattern that results when applying SPADE. In principle, since the generated spike trains are independent, no pattern should be detected in our data. An optimal result would be a false positive rate around or under the fixed significance level. Starting from this observation, we can see from the figure that the false positive rate averaged over firing rates (dashed blue line) is significantly lower than the confidence level of 5% (dashed grey line). We can interpret this may as a result of how to carry out the multiple testing. Therefore, we explain briefly how the Holm-Bonferroni correction is implemented in SPADE. The algorithm characterizes every pattern candidate contained in the data using a signature (z, c, d) (see chap. 4). Each group of patterns with the same signature is tested against the null hypothesis if its corresponding p-value is less than exactly one. A p-value of one would correspond to a pattern signature being in all surrogates.

Using the Holm-Bonferroni correction (Holm (1979)), the signature with the lowest p-value is tested against the ratio of the confidence level and the number of tests. If this test rejects the null hypothesis, the next lowest p-value is compared against the confidence level divided by the number of tests minus one. Thus, every rejection of the

¹The chosen parameter chosen for the pattern set reduction are (2, 2, 2). They are just mentioned for completeness, but will not play an important role here.




Figure 6.1: The figure shows the rate of false positives that results from SPADE analyzing a set of 20 1s long independent spike trains. In the left and right panel, we consider PPR spike trains ($\tau_r = 2.5$ ms) and gamma spike trains ($\gamma = 2$). The required p-value spectrum was not calculated using surrogates, but instead newly generated spike trains of the same type. The spike trains' underlying statistics are known, which we can use to get a ground truth null hypothesis distribution. The x -axis shows the firing rate in Hz, and the y -axis shows the false positive rate: the number of false positives averaged over the realizations of the data set. The dashed grey line represents the confidence level of 5%, while the dashed blue line shows the average of the false positive rate over firing rates.

null hypothesis lowers the number of tests by one. The first time the null hypothesis is not rejected, we stop testing further against the null hypothesis. By construction, this procedure keeps the family-wise error, i.e., the probability of having at least one false positive, below the confidence interval (Holm (1979)). It seems that in our case, the false positive rate is further decreased due to the many tests that arise from the three dimensions of the p-value spectrum. It is still an ongoing object of research in how far it is possible to decrease the number of tests, e.g., by testing for each pattern size z and duration d only the pattern signature with the highest occurrence c . Nevertheless, when comparing these results to false positive rates yielded by the usage of the surrogate methods, they shall also result in a false positive rate lower than the confidence level.

Finally, we can see that the false positive rate does not depend on the firing rate. The mean false positive rate of the gamma spike trains is higher than for PPR spike trains, but according to the standard error of the mean shown in figure 6.5, it is not possible to clearly distinguish between the two.

For the evaluation of the surrogate methods, we can say that according to the results of the ground truth null hypothesis, the false positive rates should not be much higher than 2%.

6.1.2 Null hypotheses using surrogate techniques

In order to analyze the effect of the three surrogate methods on the false positive rate of SPADE's results, we used the same datasets and the same parameters as in the previous section. The only difference here is that we used dithering surrogates (UD, UDR, JISI) for the generation of the null hypothesis. We apply SPADE to the datasets separately per dithering method, comparing our three surrogate techniques in terms of false positives (patterns resulting from each analysis).

The result can be seen in figure 6.2. In the upper plot, we observe a firing rate dependence of the false positive rate related to the uniform dithering. For the PPR spike trains of 100 Hz, this yields having an average of more than 8 false positives per realization. In contrast, the increase of the false positive rate for the gamma spike trains is smaller, but it results in a false positive rate of more than 1 for a firing rate of 100 Hz. It can be seen clearly in the figure that the first values exceeding the confidence level are for 50 Hz (PPR) and 60 Hz (gamma). We can interpret the results as a consequence of the clipped firing rate discussed in sec. 5.1 and shown especially in fig. 5.3. The difference between the spike count of the clipped spike trains and that of its clipped surrogates becomes more considerable for higher the firing rate. The effect is less intense for the gamma process as for the PPR (see fig. 5.3) and is reflected as such in the false positive rate. For firing rates until 40 Hz (PPR) or 50 Hz (gamma), the uniform dithering does not differ significantly compared to the other surrogate methods.

We can see the impact of the difference in spike count between the ground truth and uniform dithering, also comparing the p-value spectra. We show them in figure 6.3 and in the appendix (figs. C.2 - C.4). For almost all combinations of size, occurrence, and duration, the p-value obtained by uniform dithering is smaller than the ground truth p-value, independent of the spike train type. Additionally, it seems that for every size,




Figure 6.2: The figure shows the rate of false positives that results from SPADE analyzing a set of 20 1s long independent spike trains. In the left and right panel, we show PPR spike trains ($\tau_r = 2.5$ ms) and gamma spike trains ($\gamma = 2$), respectively. We calculate the p-value spectrum using surrogates created with uniform dithering (UD, orange), uniform dithering with refractoriness (UDR, green), and joint-ISI dithering (JISI, red). The x -axis shows the firing rate in Hz, and the y -axis shows the false positive rate, i.e., the number of false positives averaged over the realizations of the data set. The dashed lines represent for each surrogate method the average over firing rates. The grey dashed line shows the confidence level.

the p-value spectrum is shifted towards lower occurrences. The smaller the pattern size, the stronger is this shift. Thus, most of the false positives when using uniform dithering should arise for a small pattern size. In fact, in our data, slightly more than 50 have a pattern size of two, as can be seen in the appendix in fig. C.1. For the case shown here of a firing rate of 80 Hz and PPR spike trains, 51 of 92 false positives consist of only two spikes.

Before discussing further the results of the false positives, we first address the graphs of the p-value spectra just introduced (fig. 6.3, as well as in the appendix figs. C.2 - C.4). At a firing rate of 80 Hz, we show the p-value spectra for both spike train types and all four ways to create the spectra. To simplify the comparison between the p-value spectra generated by different methods, we also show them as contour plots (see fig. 6.4). Shifts in the p-value spectrum are, therefore, made more visible. First of all, we can see that in terms of ground truth for PPR spike trains, the p-values are generally slightly higher than for gamma spike trains. Although we do not intend to analyze this in detail, we can at least say that it can be understood as a result of the differences in the autocorrelation (Pipa et al. (2013)).

For PPR spike trains, besides the already described shift for uniform dithering, it can be noted that the p-value spectra for UDR and joint-ISI dithering are very similar to the ground truth. The closeness is because both the autocorrelation function and the clipped firing rate fit very well (see sec. 5.4).

In the case of gamma spike trains, the p-value spectra of joint-ISI dithering are very similar to the ground truth. While for the UDR, there is a comparable, smaller shift as for the uniform dithering. We attribute this to the fact that the clipped firing rate for UDR decreases in this case, like for the uniform dithering (see sec. 5.4.5).

Further analyzing figure 6.2, we can observe that for the joint-ISI dithering, the false positive rates are between 0 and 5%, with an average close to 2%. There is no observable dependence on the firing rate. Regarding the average over firing rates of the false positives (see fig. 6.5), the difference between ground truth and joint-ISI is much smaller than one standard error of the mean. Thus, even if this mean is a bit higher, it can not be stated that this is a significant tendency.

For PPR spike trains, there is mostly no difference to the usage of uniform dithering with refractoriness and joint-ISI dithering regarding the false positives as already shown in the p-value spectra (fig. C.2). In contrast, for gamma spike trains with firing rates over 80 Hz, the false positive rate increases with the firing rate for both, but more strongly for UDR. The increase is due to the change in the clipped firing rate (see sec. 5.4.5), which corresponds to the shift in the p-value spectrum.

Concluding the analysis of false positives performed here, we can affirm that the joint-ISI dithering performs very close to the ground truth null hypothesis. While the uniform dithering shows a substantial increase in false positives for high firing rates, and the UDR also shows this increase, it only starts at very high firing rates.

In further research, this analysis of false positives needs to be extended not only to varying rates across neurons but also to non-stationary rate profiles for single neurons.




Figure 6.3: The figure shows the p-value spectrum corresponding to 20 independently generated PPR ($\tau_r = 2.5$ ms) spike trains of 1 s (upper panel) and to uniformly dithered spike trains (lower panel). In both cases, it is an average over the 100 data sets used to analyze the false positives of this specific firing rate. The three dimensions of the p-value spectrum are shown as such that every size has one subpanel. On the x -axis are the occurrences and on the y -axis the durations in bins. We show the p-values according to the logarithmic colormap on the right side.




Figure 6.4: The figure shows contour lines for the p-value spectra for PPR spike trains (upper panel) and gamma spike trains (lower panel). The p-value spectra themselves are shown in figures 6.3 & C.2 - C.4. The contour lines correspond to p-values of 50% (marked with squares), 5% (pluses), 0.5% (circles) and 0.05% (xs). We calculate the exact marker positions employing linear interpolation. The four ways to create the p-value spectra we present in different colors: ground truth null hypothesis (GT, blue), uniform dithering (UD, orange), joint-ISI dithering (JISI, red) and UDr (green). The three dimensions of the p-value spectrum are shown as such that every size has one subpanel. On the x -axis are the occurrences and on the y -axis the durations in bins.




Figure 6.5: The figure shows the average of the false positives rates displayed in figures 6.1 & 6.2 over the firing rate. In addition to the mean, the error bars represent the standard error of the mean. It is shown for PPR (left) and gamma (right) spike trains with the null hypotheses: ground truth (GT, blue), uniform dithering (UD, orange), uniform dithering with refractoriness (UDR, green) and joint-ISI dithering (JISI, red). Since the average false positive rate of the uniform dithering is high (see fig. 6.2), it is not shown in the panel for the PPR spike trains passing the limits of the scale.

6.2 False negatives

The analysis of false negatives is closely related to how we think the characteristics of true positives could be. In particular, one should consider whether the statistics of spike trains with spatio-temporal patterns should be different from those of the other spike trains without patterns. Furthermore, one may ask to what extent the occurrences of true positives should stand out from those of by-chance patterns. We have limited ourselves here to generating spike trains with spatio-temporal patterns following the same ISI distributions as the independent spike trains. This generation is more subtle for PPR and gamma spike trains than for Poisson spike trains, which were analyzed in Quaglio et al. (2017); Stella et al. (2019). False negatives in this setting would be then: patterns suitably injected into artificial data, and not detected by SPADE. By comparing the number of false negatives across the different surrogate techniques, we can understand which methods are more appropriate to detect patterns in electrophysiological data. In order to choose the number of occurrences of the true positive patterns appropriately, we have taken the approach to select them according to the p-value spectrum of independent spike trains. The results expressed in terms of false negatives finally make it possible to distinguish which surrogate method is most likely to detect true positives.

6.2.1 Data generation with spatio-temporal patterns

In order to create PPR and gamma spike trains with spatio-temporal patterns, we use two properties of Poisson spike trains. First, the superposition of Poisson spike trains is again a Poisson spike train. This property allows us to generate a set of spike trains with a spatio-temporal pattern as a superposition of a shared pattern spike train with firing rate λ_p and an individually generated background spike train with reduced firing rate $\lambda_r = \lambda - \lambda_p$ (Cardanobile and Rotter (2010)). Secondly, it is possible to generate PPR and gamma spike trains (if the shape factor γ is an integer) starting with Poisson spike trains. The Poisson process with refractoriness is achieved from a Poisson process by deleting every spike that is closer to its preceding spike than the refractory period τ_r (see appendix C.3). To delete every but the γ -th spike effects the transition from the Poisson to the gamma process instead (Heeger (2000)). So that the resulting PPR or gamma spike train has a firing rate equal to λ , the firing rate of the associated Poisson spike train needs to be $\tilde{\lambda}(\lambda) = \lambda/(1 - \lambda\tau_r)$ and $\tilde{\lambda}(\lambda) = \gamma\lambda$, respectively.

Combining this, we can create a spatio-temporal pattern containing n spikes by creating one Poisson spike train representing the first spike of the patterns, and by copying it n times. In this way, we create spike trains representing the patterns only, where each spike train has a rate of $\tilde{\lambda}(\lambda_p)$. In order to generate patterns with a particular configuration of delays between spikes, we shift the "pattern spike trains" appropriately by the delay. Moreover, for each of the n spike trains, we create another Poisson spike train of rate $\tilde{\lambda}(\lambda - \lambda_p)$, representing the background activity. Pattern spike trains and background spike trains are then superimposed, leading to the final version of the spike train, containing both the inserted pattern and additional independent spikes. Spikes are deleted in this superposition of spike trains (differently for PPR and gamma processes)

in the manner described above.

Our data sets contain 100 realizations (as for the analysis of false positives). Moreover, we decide to shift only the last of the n spike trains to obtain the duration, e.g., a pattern of three spikes and three bins long will have the first two spikes synchronous, and the last spike will occur two bins later. In each realization, we insert only one pattern. Finally, for each combination of duration and size, patterns are equally represented in the number of realizations. Sizes vary from 2 to 5 spikes, while durations are fixed in the range of 1 to 5 bins.

The occurrences for each of these inserted patterns are chosen such that in the p-value spectrum of independent spike trains (see figs. 6.3 & C.3), it would be the minimum number of occurrences for which the corresponding p-value is below 0.05%. With this choice, it is to expect that even with a high number of multiple tests of 100, SPADE should detect these patterns and identify them as significant.

We call $c_{z,d}(\lambda)$ the fixed number of occurrences of each generated pattern. Since we operatively generate pattern spike trains with a rate λ_p , different realizations do not generally lead us to this $c_{z,d}(\lambda)$. However, there is a relationship between the pattern rate λ_p and the expected value of the occurrences $\bar{c}_z(\lambda_p; \lambda)$, which we determine numerically². We do this for each combination of firing rate λ and pattern size z by generating this set of spike trains for different λ_p and averaging the resulting values of $c_z(\lambda_p; \lambda)$ for each λ_p . The intermediate values we estimate via linear interpolation. In order now to create a pattern of occurrences $c_{z,d}(\lambda)$, we set the pattern rate according to $\bar{c}_z(\lambda_p; \lambda) = c_{z,d}(\lambda)$ and generate this set of spike trains until the occurrence of the inserted pattern corresponds precisely to the value $c_{z,d}(\lambda)$.

With this procedure, we generate data for the same values of firing rates as in the previous section, and we analyze this data with SPADE.

6.2.2 Analysis of data with spatio-temporal patterns

The analysis of the data with the true positive spatio-temporal patterns was mainly with the same parameters as the analysis for false positives in sec. 6.1. The only difference is in the parameters of the pattern set reduction. The pattern set reduction tests two patterns against each other that contain a set of identical spikes. We set the parameters such that SPADE identifies more likely the pattern of the larger size as significant³.

Before measuring the number of false negatives, it is necessary to define a false negative. In principle, a false negative is a pattern injected in the data and not detected by SPADE. We relax this definition by fixing that a not detected, injected pattern will not be classified as a false negative, if it is overlapping within at least two neurons of one of the patterns detected by SPADE (and conserving the lags between the spikes). Accordingly, we define a false positive as any pattern detected by SPADE for which neurons and lag configurations do not overlap in the described manner with the injected

²We do not distinguish here for the different durations since the differences are negligible as being only a cause of the boundaries. Thus, we determine $\bar{c}_z(\lambda_p; \lambda)$ only for synchrony patterns.

³The corresponding parameters are (1,3,2).

pattern. Finally, we calculate false negative/false positive rates by averaging the number of false negatives/false positives over the data sets.

We show the resulting false negative rates in figure 6.6. The first observation is that by using uniform dithering, there are slightly less false negatives than if the ground truth p-value spectrum is applied. For the data sets of gamma spike trains, we detect no false negatives when using uniform dithering. We can relate this result to what we discussed in the false-positive section (see sec. 6.1.2): Due to the shift in the p-value spectrum, it is more likely to detect false positives but also true positives

We observe that we have higher false negative rates for both uniform dithering with refractoriness and joint-ISI dithering than by applying the ground truth null hypothesis. For PPR spike trains, we further observe a firing rate dependence: false negative rates increase for firing rates higher than 60 Hz. Again, this can be related to what we have already discussed in the sections 5.4.2 and 5.4.5. On the one hand, we have seen that there is a probability that not enough spikes are displaced, possibly causing not to destroy the patterns when creating surrogates. On the other hand, for UDR, there is a very close match in terms of the clipped firing rate. Thus, since we do not have a shift in the p-value spectrum, injected patterns not destroyed in the surrogates (even if it is only a tiny number) can yield them to be not classified as significant. For the joint-ISI dithering and PPR spike trains, it is a similar situation, but the clipped firing rate is slightly lower than for the UDR (see sec. 5.4.5). The resulting shift in the p-value spectrum then compensates for destroying less spatio-temporal patterns, when creating the surrogates.

For gamma processes, the situation is reversed. Joint-ISI dithering yields a higher false negative rate in comparison to the UDR. Again, we explain this by the effective displacement and the clipped firing rate (see secs. 5.4.2 & 5.4.5). For the UDR, the clipped firing rate of the surrogates decreases more strongly than for the data itself. As a result, the p-value spectrum of the UDR is shifted (see it for the independent data in fig. C.4). Thus, too small p-values are assigned to the patterns' signatures. Consequently, to identify a pattern as significant, fewer occurrences are necessary, and the false negative rate is small.

For the joint-ISI dithering this is not the case, where there is no shift in the p-value spectrum (see fig. C.4). Since the clipped firing rate of the joint-ISI dithered spike trains follow closely that of the artificial spike trains (see sec. 5.4.5). Therefore, the case that of not destroyed spatio-temporal patterns is more relevant, which can happen if spikes stay in the same bin creating the surrogates (see sec. 5.4.2). However, this can also be due to a poor estimate of the joint-ISI distribution if there are only a few spikes (see secs. 2.3 & 5.3). Correspondingly we observe that, on average, the false negative rate is higher than all other surrogates (see 6.6).

Finally, it is worth noting that in all considered cases, the average false negative rate is below 5%. Thus, it seems appropriate to evaluate the three surrogate methods with higher weighting concerning false positives, which have larger values than the false negatives (see sec. 6.1.2).

One still may ask the question: To what extent does the presence of true positives




Figure 6.6: The upper figure shows the rate of false negatives that results from SPADE analyzing a set of 20 1 s long spike trains with injected spatio-temporal patterns. We consider PPR spike trains ($\tau_r = 2.5$ ms, left panel) and gamma spike trains ($\gamma = 2$, right panel). The required p-value spectrum was calculated with the ground truth null hypothesis (GT, blue), or using surrogate methods: uniform dithering (UD, orange), UDR (green), and joint-ISI dithering (JISI, red). The x -axis shows the firing rate in Hz, and the y -axis shows the false negative rate as defined above. The average of the false negative rates over firing rates is shown in dashed lines. The lower figure shows the same averages, but additionally, the standard errors of the mean as error bars.

impact the number of false positives? To investigate this, we have also evaluated the false positives for these data sets, which we show in figure 6.7. We find that there is not much difference between the false positive rates of these data sets with patterns and those of the independent spike train data sets (see figs. 6.1 & 6.2). The only observed differences are higher false positive rates for uniform dithering analyzing PPR spike trains at firing rates of 30 and 40 Hz. We conclude that the maximum firing rate up to which uniform dithering is reliable regarding false positives may be lower than the 40 Hz given in the previous section.

Many statistical quantities are used that combine false positives and false negatives. We utilize the false discovery rate here, i.e., the ratio of false positives over the patterns identified as significant. Thus, for any pattern classified as significant, the probability can be indicated that it is a false positive. The results for this, we show in figure 6.8. For most values, the false discovery rate is very similar to the false positive rate, even if the scale is different (the false discovery rate cannot be higher than 1). In order to obtain more relevant information from combined statistics' measures, we would need to analyze data sets with further inserted spatio-temporal patterns.

As a summary of our analysis of false positives, we can firstly state that the choice of a suitable surrogate method depends on the firing rate. The uniform dithering method performs well for low firing rates concerning either false positives or false negatives. However, for higher firing rates, the number of false positives increases dramatically. As a consequence, we do not recommend uniform dithering for spike train data that are either regular or have a refractory period to avoid possible sources of false positives.

In the comparison between UDR and joint-ISI dithering, there are few differences in terms of false positives and false negatives. The results of false positives are very similar except for an increase above a firing rate of 90 Hz for UDR used to analyze gamma spike trains. Except for this case, the false positive rate is well below the confidence level for both methods. As far as false negatives are concerned, UDR shows generally better results for false negatives. Only above 70 Hz in the analysis of PPR spike trains the results are slightly higher. Therefore, we recommend using uniform dithering with refractoriness in the analysis with SPADE of spike train data with firing rates below 70 Hz, which are regular or have a refractory period. As a reference point, we recommend choosing the neuron with the highest firing rate. The joint-ISI dithering can be used very well over the complete range of firing rates. For high firing rates, it is the only choice we recommend, robust in both false positives and false negatives.




Figure 6.7: The figure shows the rate of false positives that results from SPADE analyzing the data sets with the inserted true positives. In the left and right panel, we show PPR spike trains ($\tau_r = 2.5$ ms) and gamma spike trains ($\gamma = 2$), respectively. The p-value spectrum was calculated using independent spike trains (GT, blue) and surrogates created with uniform dithering (UD, orange), uniform dithering with refractoriness (UDR, green) and joint-ISI dithering (JISI, red). The x -axis shows the firing rate in Hz, and the y -axis shows the false positive rate, i.e., the number of false positives averaged over the realizations of the data set. The dashed lines represent for each surrogate method the average over firing rates. The grey dashed line shows the confidence level.




Figure 6.8: The figure shows the false discovery rate that results from SPADE analyzing the data sets with the inserted true positives, i.e., the ratio of the number of false positives over the number of patterns detected as positives. In the left and right panel, we show PPR spike trains ($\tau_r = 2.5$ ms) and gamma spike trains ($\gamma = 2$), respectively. We calculate the p-value spectrum using independent spike trains (GT, blue) and surrogates created with uniform dithering (UD, orange), uniform dithering with refractoriness (UDR, green) and joint-ISI dithering (JISI, red). The x -axis shows the firing rate in Hz, and the y -axis shows the false discovery rate. The dashed lines represent for each surrogate method the average over firing rates.

Chapter 7

Conclusion

7.1 Discussion

This thesis aimed to compare different dithering methods concerning the creation of a statistically robust null hypothesis for the detection of significant spatio-temporal spike patterns. Therefore, we not only provided an introduction to the world of neurons and spikes but also summarized the basics of statistical point processes used to describe spike trains. We have decided to focus on Poisson processes with refractoriness (PPR) and gamma processes, allowing us to include an absolute or relative refractory period. Thus, they represent two ways to model regular spike trains¹. It is worth noting that we have extended the analysis of PPR spike trains, carried out by Deger et al. (2012), also to include the basics of non-stationary PPR spike trains. By analyzing the characteristics of electrophysiological recordings, exemplified by the reach-to-grasp data set, we ensured that it is legitimate to restrict ourselves to regular spike trains. However, in this thesis, we did not consider the property of recorded neurons that firing rates change across neurons and with time but limit ourselves to firing rates that are stationary and the same for all neurons.

By introducing SPADE as a method for spatio-temporal pattern detection, we showed how patterns are obtained from a spike train data set, and illustrated how they are tested for significance. The null hypothesis is that all spatio-temporal patterns are only a by-product of the statistical properties of mutually independent spike trains. In order to approximate these independent spike trains, SPADE uses surrogates. The surrogate method should, therefore, ensure independence between the spike trains while maintaining the statistical properties. These surrogates are then used to generate a distribution of p-values called the p-value spectrum. With this objective, we have compared three dithering methods. These are uniform dithering (a standard approach), uniform dithering with refractoriness (UDR - a newly introduced adaptation of the uniform dithering), and joint-ISI dithering (approach explicitly seeking to preserve the inter-spike interval (ISI) distribution). Note that we defined the joint-ISI dithering somewhat differently

¹Thereby, we have limited ourselves to gamma processes with shape factors $\gamma > 1$.

than Gerstein (2004) and Louis et al. (2010b) and discussed the reasons for this in detail.

We have identified the interplay between the *method of time discretization* (clipping) and the surrogates as a sticking point in the use of surrogate methods within SPADE. Clipping, the first step before applying the pattern mining algorithm, is an essential component of SPADE. We pointed out that the drifting apart between the firing rate of clipped spike trains and that of clipped surrogate spike trains is the primary source of the false positives that result from uniform dithering. For this dithering version, we presented the clipped firing analytically and numerically. In doing so, we have observed that *the clipped firing rate decreases drastically* in the case of high underlying firing rates.

Also, for the UDR and joint-ISI, we have investigated this numerically. As a result, we observed that the UDR is very robust with respect to the effect of clipping when analyzing PPR spike trains, but behaves similar to the uniform dithering when analyzing gamma spike trains. However, the joint-ISI dithering shows only small variations in firing rate after clipping regardless of the type of spike train. We also showed the clipped spike count that emerges from analyzing the reach-to-grasp data set with the surrogate techniques. Thereby, we found no substantial differences between the UDR and the joint-ISI dithering. For the uniform dithering, however, we could see a marked reduction of the clipped spike count.

Moreover, we studied the autocorrelation of the surrogate spike trains and their ISI distribution. We showed that both measures are not conserved within a uniform dithered surrogate spike train. In contrast, UDR and joint-ISI dithering preserve the autocorrelation and ISI distribution of PPR spike trains approximately. When applied to gamma spike trains, only joint-ISI dithering is able to preserve these measures.

Additionally, we studied the differences of the underlying spike trains to their respective surrogates in terms of cross-correlation and the number of spikes that did not change their position in bins during the transition to the surrogate. We observed first that for all three dithering methods, the number of spikes in equal bins in the surrogates and the original spike trains is higher than when comparing two independent spike trains. In principle, this can lead to patterns not being sufficiently destroyed. Particularly marked is this rate of spikes in the same bins when applying UDR to gamma spike trains.

Moreover, we could shed some light on the extent to which serial ISI correlations emerge in the surrogate spike trains. While uniform dithering causes negative serial correlations, joint-ISI dithering causes positive serial correlations, but for the UDR, it depends on the spike train model. For PPR spike trains, it shows positive serial correlation, whereas, for gamma spike trains, we observe negative serial correlations.

After this comparison of the surrogate methods in terms of single spike train statistics, we designed and generated a test data set. As models for the generation of spike trains, we decided to use PPR as well as Gamma spike trains. The artificial data sets comprised both independent data and spike trains with injected patterns. On these data sets, we applied then SPADE with each of the surrogate methods. In this way, we have evaluated how far the null hypotheses generated by the three different surrogate techniques were adequate in terms of false positives and false negatives.

For the case of independent spike train data, we were able to provide not only an estimate of the p-value spectrum but also the ground truth as a reference for the comparison. This ground truth resulted from the known, underlying statistics of the generated PPR and gamma spike trains. Regarding uniform dithering, we observed the already predicted *firing rate dependent increase of false positives*. For the other two methods, no significant difference to the ground truth could be detected, except for the analysis of gamma spike trains using UDR at high firing rates. There, we see a firing rate-dependent increase in false positives, but not as strong as in the case of uniform dithering.

In the second step of the data sets with inserted spatio-temporal patterns, it was additionally possible to evaluate the surrogate methods regarding false negatives. Here, we observed that uniform dithering scored better than the other methods. However, this could not outweigh the poor performance on the false positives. The false negative rate resulting from UDR or joint-ISI dithering was generally below 5 %. The only deviation we found was in the analysis of PPR spike trains using UDR, where we recognized a slight increase of false negatives for higher firing rates.

Louis et al. (2010b) have already compared uniform dithering and joint-ISI dithering (and other methods). In the test case carried out there, the same tendencies were visible that there are more false positives for uniform dithering and more false negatives for joint-ISI dithering. However, the vast increase of false positives in uniform dithering, observed in this thesis, could not be observed by Louis et al. (2010b) using a different way to discretize time. Indeed, this increase in false positives is a clipping specific problem.

We concluded that we recommend using *uniform dithering with refractoriness* for spike train data sets where the highest firing rate is 70 Hz maximum. However, if a data set contains neurons with higher firing rates, we recommend applying the *joint-ISI dithering*.

The results and the procedure presented in this thesis yield a recommendation on which surrogate method to employ within SPADE. Nonetheless, we also would like to mention the relevant aspects that we have not covered. First of all, spike trains in recorded data sets have non-stationary firing rate profiles and vary strongly across neurons. We could not analyze this within the scope of this thesis, but a logical next step includes expanding our evaluations to these cases. For this purpose, we already started analyzing independent data with firing rate profiles corresponding to those of the reach-to-grasp data set. This analysis has not yet concluded, even though a few preliminary results are already available.

In the analysis of non-stationary spike train data, another issue would be whether the dithering methods should make use of operational time (see sec. 2.4.3.1), as was done in Louis et al. (2010b). The use of operational time has advantages and disadvantages. In the test case discussed in that paper, the methods can achieve even fewer false positives by performing the dithering in operational time. Crucial for us, not to use this, are the uncertainties associated with the estimation of the firing rate (see sec. 2.3), especially since we do not necessarily want to make assumptions about how the firing rate changes over trials.

Another important aspect not mentioned is that there are also surrogate methods

that do not rely on dithering. Indeed, Louis et al. (2010a) provide a comprehensive overview of surrogate methods. For example, it would be possible to create surrogates of already discretized spike trains and thus to avoid the problems caused by clipping. A currently discussed approach to do so is called bin shuffling, where the bin contents are shuffled within a sliding window.

It would also be possible not to displace each spike individually, but to move the entire spike train on a trial-by-trial basis (Louis et al. (2010a,b)). This approach is straightforward and, in any case, maintains the spike train’s statistical properties. It remains to analyze the extent to which creating these surrogates destroys the existing patterns, and if the change in firing rate correlations would have an impact. In terms of coincidence detection, however, Louis et al. (2010a) have shown that this shifting approach results in a low number of false negatives but a high number of false positives. Whether this would also be the case when using SPADE can hardly be predicted.

Recently there have been some discussions on the subject of multiple testing that we discussed in sec. 6.1.1. The idea is to calculate the number of tests differently. As a consequence, the false positive rate based on the ground truth null hypothesis should approach closer to the confidence level.

In this thesis, we used Poisson processes with refractoriness and gamma processes as spike train models. These are not the only renewal processes used to model spike trains. Also widely used are the Log-Normal processes and the Inverse Gaussian processes (van Vreeswijk (2010)), which are somewhat more complex than those presented here. Nevertheless, we think that the extension to non-stationary firing rates is a more decisive issue, given that the underlying research question is to evaluate electrophysiological data.

Torre et al. (2016) already examined the reach-to-grasp data for zero-lag patterns with an earlier version of SPADE. Therefore, it is only natural to ask whether our results impact these analyses. However, this paper used a different mining algorithm (Torre et al. (2016)), which makes it difficult to predict if there is any impact. This mining algorithm was not bin-centered but spike-centered, so the spike count was not reduced by clipping. Unfortunately, this approach is feasible only for the case of synchronous patterns.

Nonetheless, we have already started to analyze the reach-to-grasp data set with the newest version of SPADE using different surrogate techniques to see whether different results arise. These analyses should not only detect spatio-temporal patterns and make it possible to relate these patterns to different behavioral contexts. Indeed, this is the ultimate purpose of the development of SPADE.

In this development, there had already been the incorporation of time delays (Quaglio et al. (2017)) and of a corresponding third dimension p-value spectrum (Stella et al. (2019)). With the study of this master thesis, the entire methodological work shall now lead to more substantiated results of the analysis of experimental records.

7.2 Implementation

We added some functions to the open-source Python package `elephant` during the work on this master's thesis. First, these are the functions for generating stationary and non-stationary Poisson processes with refractoriness. Secondly, we added the uniform dithering with refractoriness as a possible adaptation to the uniform dithering function. Finally, the joint-ISI dithering was newly implemented, roughly following the pseudocode in Louis et al. (2010a). Additionally, we optimized the implementation of SPADE for run time. Here we show an example script using these functions.

```
1 import elephant.spike_train_generation as gen
2 import elephant.spike_train_surrogates as surr
3 from elephant.spade import spade
4 import quantities as pq # Library to use units
5 import neo # Library to work with electrophysiological data
6
7 refractory_period = 2.5 * pq.ms
8 # Generation of stationary PPR spike train
9 stationary_spiketrain = gen.homogeneous_poisson_process(
10     rate=80. * pq.Hz, t_start=0. * pq.s, t_stop=1. * pq.s,
11     refractory_period=refractory_period)
12
13 # Generation of non-stationary PPR spike train
14 # with a step-function as rate profile
15 firing_rate_profile = neo.AnalogSignal(
16     signal=[10., 80.], units=pq.Hz, sampling_period=0.5*pq.s)
17 non_stationary_spiketrain = gen.inhomogeneous_poisson_process(
18     rate=firing_rate_profile, refractory_period=refractory_period)
19
20 # Creation of surrogates using the uniform dithering with refractoriness
21 # and using the joint-ISI dithering.
22 dither_parameter = 15. * pq.ms
23 for spiketrain in (stationary_spiketrain, non_stationary_spiketrain):
24     surrogate_udr = surr.dither_spikes(
25         spiketrain=spiketrain, dither=dither_parameter,
26         refractory_period=refractory_period)[0]
27     surrogate_jisi = surr.JointISI(
28         spiketrain=spiketrain, dither=dither_parameter).dithering()[0]
29
30 # Application of SPADE
31 spade_results = spade(
32     spiketrains=[stationary_spiketrain, non_stationary_spiketrain],
33     binsize=3*pq.ms, winlen=5, min_spikes=2, min_occ=10, min_neu=2,
34     n_surr=5000, alpha=0.05, stat_corr='holm', spectrum='3d#',
35     surr_method='joint_isi_dithering', psr_param=[1,3,2])
```

Example usage of the Python functions

For the sake of reproducibility, we show all program code needed to recreate the results of this thesis in the GitHub repository:

https://github.com/INM-6/dithering.methods_master_thesis.

7.3 Acknowledgements

First of all, I would like to thank Prof. Sonja Grün, who gave me the opportunity to do this master's thesis at the INM-6 at the research center Jülich. I really appreciated the guidance, the provision of ever new ideas, and the exciting, often very lively discussions.

Secondly, I would like to thank my supervisor Alessandra Stella. I would especially like to thank her for this close supervision. The cooperation in this closely interwoven project was a great pleasure for me. All the time that has been devoted to me, the effort to be familiar with all the details and the support and advice that I have received have become very important to me. It was particularly gratifying for me that every line of this thesis was studied and commented on when necessary.

I would also like to thank Prof. Riccardo Mazzarello, who has agreed to act as a reviewer.

Furthermore, I would like to thank Prof. Günther Palm and Prof. Alexa Riehle for fruitful discussions.

Special thanks to Danylo Ulianych for his help in adding functions to the elephant package.

In addition, I would like to thank Dr. Desy Askitosari and Marvin Schossau for proofreading.

Finally, I would like to thank all colleagues at the INM-6, who all contribute to an excellent working atmosphere.

Appendix A

Calculations for Point Processes

A.1 Equivalence of firing rate definitions

We show here that the two definitions of the firing rate in equations (1.1) and (2.5) are equivalent. We start with the first one:

$$\lambda(t) = \lim_{\epsilon \rightarrow 0} \frac{1}{\epsilon} \text{Prob}(\text{spike in } [t, t + \epsilon]) \quad \text{for } t \geq 0, \epsilon > 0.$$

If we only look at the probability term, we can transform it as:

$$\begin{aligned} \text{Prob}(\text{spike in } [t, t + \epsilon]) &= 1 - \text{Prob}(\text{no spike in } [t, t + \epsilon]) \\ &= 1 - \prod_{i=1}^{\infty} \text{Prob}(t_i \text{ not in } [t, t + \epsilon]) \\ &= 1 - \prod_{i=1}^{\infty} (1 - \langle \Theta(t_i - t) \Theta(t + \epsilon - t_i) \rangle_s + \mathcal{O}(\epsilon^2)) \\ &= 1 - \prod_{i=1}^{\infty} (1 - \langle \delta(t_i - t) \rangle_s \epsilon + \mathcal{O}(\epsilon^2)) \\ &= \sum_{i=1}^{\infty} \langle \delta(t_i - t) \rangle_s \epsilon + \mathcal{O}(\epsilon^2) \end{aligned}$$

In the third line, we have used that the probability of having a spike in an interval $[t, t + \epsilon]$ is up to terms of order $\mathcal{O}(\epsilon^2)$ equal to the expected number of spikes in $[t, t + \epsilon]$.

Plugging this into the upper firing rate equation yields:

$$\begin{aligned}
\lambda(t) &= \lim_{\epsilon \rightarrow 0} \frac{1}{\epsilon} \text{Prob}(\text{spike in } [t, t + \epsilon]) \\
&= \lim_{\epsilon \rightarrow 0} \frac{1}{\epsilon} \left(\sum_{i=1}^{\infty} \langle \delta(t_i - t) \rangle_s \epsilon + \mathcal{O}(\epsilon^2) \right) \\
&= \sum_{i=1}^{\infty} \langle \delta(t_i - t) \rangle_s \stackrel{(2.4)}{=} \langle s(t) \rangle_s
\end{aligned}$$

Thus, we have shown that both definitions are equivalent. Note that all steps performed here can also be done in the other direction.

A.2 Renewal processes

We show here the relationship between the firing rate, the distribution of the first spike and the positive- τ conditional firing rate:

$$\begin{aligned}
\lambda(t) &\stackrel{(2.5, 2.4)}{=} \left\langle \sum_{i=1}^{\infty} \delta(t - t_i) \right\rangle_s \tag{A.1} \\
&\stackrel{(2.3)}{=} \sum_{i=1}^{\infty} \int \prod_{j=0}^i dt_j \rho_j(t_j | t_{j-1}, \dots, t_1) \delta(t - t_i) \\
&\stackrel{(2.30)}{=} \sum_{i=1}^{\infty} \int dt_i (\rho_1 \circledast \underbrace{p \circledast \dots \circledast p}_{(i-1) \text{ times}})(t_i) \delta(t - t_i) \\
&= \rho_1(t) + \left(\rho_1 \circledast \sum_{i=2}^{\infty} p^{(i-1) \circledast} \right)(t) \\
&\stackrel{(2.33)}{=} \rho_1(t) + (\rho_1 \circledast \lambda_c^+)(t)
\end{aligned}$$

From the second to the third line, we could integrate out over all t_j with $j > i$.

Further, we analyze the relationship between the autocorrelation function and the conditional firing rate:

$$\begin{aligned}
\psi(t, t') &= \left\langle \sum_{i=1}^{\infty} \sum_{j=1}^{\infty} \delta(t - t_i) \delta(t' - t_j) \right\rangle_s \\
&= \left\langle \sum_{i=1}^{\infty} \sum_{j=i+1}^{\infty} \delta(t - t_i) \delta(t' - t_j) \right\rangle_s \\
&\quad + \left\langle \sum_{i=1}^{\infty} \delta(t - t_i) \delta(t' - t_i) \right\rangle_s \\
&\quad + \left\langle \sum_{i=1}^{\infty} \sum_{j=1}^{i-1} \delta(t - t_i) \delta(t' - t_j) \right\rangle_s.
\end{aligned} \tag{A.2}$$

Calculating the different summands by its own and using the symmetry between summand 1 and 3, only the following two terms need to be considered.

$$\begin{aligned}
&\left\langle \sum_{i=0}^{\infty} \sum_{j=i+1}^{\infty} \delta(t - t_i) \delta(t' - t_j) \right\rangle_s \\
&\stackrel{(2.30)}{=} \sum_{i=1}^{\infty} \sum_{j=i+1}^{\infty} \int \prod_{k=0}^i dt_k \rho_k(t_k | t_{k-1}, \dots, t_1) \times \\
&\quad \times \int \prod_{l=i+1}^j dt_l \rho_l(t_l | t_{l-1}, \dots, t_1) \delta(t - t_i) \delta(t' - t_j) \\
&\stackrel{(2.29)}{=} \sum_{i=1}^{\infty} \sum_{j=i+1}^{\infty} \int dt_i (\rho_1 \otimes \underbrace{p \otimes \dots \otimes p}_{(i-1) \text{ times}})(t_i) \times \\
&\quad \times \int dt_j (\underbrace{p \otimes \dots \otimes p}_{(j-i) \text{ times}})(t_j - t_i) \delta(t - t_i) \delta(t' - t_j) \\
&= \sum_{i=1}^{\infty} \sum_{j=i+1}^{\infty} (\rho_1 \otimes p^{(i-1) \otimes})(t) p^{(j-i) \otimes}(t' - t) \\
&\stackrel{(2.33, 2.35)}{=} \lambda(t) \sum_{k=1}^{\infty} p^{(k) \otimes}(t' - t)
\end{aligned} \tag{A.3}$$

$$\left\langle \sum_{i=0}^{\infty} \delta(t - t_i) \delta(t' - t_i) \right\rangle = \lambda(t) \delta(t - t') \tag{A.4}$$

This can be summarized by using the conditional firing rate (see eq. (2.33)).

$$\psi(t, t') = \lambda(t) \sum_{k=1}^{\infty} p^{(k)\otimes}(t' - t) + \lambda(t) \delta(t - t') + \lambda(t') \sum_{k=1}^{\infty} p^{(k)\otimes}(t - t') \quad (\text{A.5})$$

$$\stackrel{(2.33)}{=} \lambda(t) \delta(t - t') + \lambda(t) \lambda_c^+(t' - t) + \lambda(t') \lambda_c^+(t - t') \quad (\text{A.6})$$

$$\stackrel{(2.34)}{=} \lambda \delta(t - t') + \lambda \lambda_c(|t - t'|) \quad (\text{A.7})$$

The fact that the position of the first spike follows eq. 2.36 yields a close formula for the p.d.f of the first spike. To make it more explicit that eq. 2.36 is only defined for $t \geq 0$, we replace on the left side $\lambda \rightarrow \lambda \Theta(t)$. Therefore, it follows that:

$$(\rho_1 \otimes \lambda_c)(t) + \rho_1(t) = \lambda \Theta(t).$$

Now convolving this equation with $p(\tau)$ yield:

$$\begin{aligned} \Rightarrow \underbrace{(p \otimes \rho_1 \otimes \lambda_c)(t) + (p \otimes \rho_1)(t)}_{\stackrel{2.33}{=} (\rho_1 \otimes \lambda_c)(t)} &= \lambda \int_0^t d\tau p(\tau), \\ \Leftrightarrow \lambda \Theta(t) - \rho_1(t) &= \lambda \int_0^t d\tau p(\tau). \end{aligned} \quad (\text{A.8})$$

A.3 Markov processes

To calculate the ISI distribution for Markov processes, it is necessary to plug eq. (2.41) into eq. (2.16).

$$f(\tau) = \lim_{T \rightarrow \infty} c_1(T) \left\langle \sum_{i=1}^{\infty} \delta(\tau - (t_{i+1} - t_i)) \Theta(T - t_{i+1}) \right\rangle_s$$

The second term becomes

$$\begin{aligned}
& \left\langle \sum_{i=1}^{\infty} \delta(\tau - (t_{i+1} - t_i)) \Theta(T - t_{i+1}) \right\rangle_s \\
&= \sum_{i=1}^{\infty} \int \left(\prod_{j=1}^{i+1} dt_j \right) \left(\prod_{j=2}^{i+1} p(t_j | t_{j-1}) \right) \rho_j(t_1) \delta(\tau - (t_{i+1} - t_i)) \Theta(T - t_{i+1}) \\
&\stackrel{(2.48)}{=} \sum_{i=1}^{\infty} \int_0^{\infty} dt_i \int_0^T dt_{i+1} \rho_i(t_i) p(t_{i+1} | t_i) \delta(\tau - (t_{i+1} - t_i)) \\
&\stackrel{(2.49)}{=} \int_0^{\infty} dt \lambda(t) \int_0^T dt' p(t' | t) \delta(\tau - (t' - t)) \\
&= \int_{\tau}^T dt' \lambda(t' - \tau) p(t' | t' - \tau) \\
&= \int_0^{T-\tau} dt \lambda(t) p(t + \tau | t)
\end{aligned}$$

The normalization factor is

$$\begin{aligned}
c_1(T) &= \frac{1}{E(N(T)) - \langle \Theta(T - t_1) \rangle_s} \\
&= \frac{1}{\int_0^T dt \lambda(t) - \int_0^T dt \rho_1(t)}.
\end{aligned}$$

The ISI distribution is thus expressed as

$$\begin{aligned}
f(\tau) &= \lim_{T \rightarrow \infty} \frac{\int_0^{T-\tau} dt \lambda(t) p(t + \tau | t)}{\int_0^T dt \lambda(t) - \int_0^T dt \rho_1(t)} \\
&= \lim_{T \rightarrow \infty} \frac{\int_0^T dt \lambda(t) p(t + \tau | t)}{\int_0^T dt \lambda(t)}.
\end{aligned}$$

Following the same steps the Joint-ISI distribution becomes

$$\begin{aligned}
f(\tau, \tau') &= \lim_{T \rightarrow \infty} \frac{\int_0^{T-(\tau+\tau')} dt \lambda(t) p(t + (\tau + \tau') | t + \tau) p(t + \tau | t)}{\int_0^T dt \lambda(t) - \langle \Theta(T - t_2) \rangle_s - \langle \Theta(T - t_1) \rangle_s} \\
&= \lim_{T \rightarrow \infty} \frac{\int_0^T dt \lambda(t) p(t + (\tau + \tau') | t + \tau) p(t + \tau | t)}{\int_0^T dt \lambda(t)}.
\end{aligned}$$

A.4 Poisson process with refractoriness

In order to simplify the expression of the C_{V2} for Poisson processes with refractoriness, we apply a variable transformation.

$$\begin{aligned}
C_{V2} &= 2 \left\langle \frac{|\tau - \tau'|}{\tau + \tau'} \right\rangle_{\tau, \tau'} \\
&\stackrel{(2.32)}{=} \int_0^\infty d\tau \int_0^\infty d\tau' p(\tau) p(\tau') \frac{|\tau - \tau'|}{\tau + \tau'} \\
&= 2(\tilde{\lambda})^2 \int_{\tau_r}^\infty d\tau \int_{\tau_r}^\infty d\tau' \exp(-\tilde{\lambda}(\tau + \tau' - 2\tau_r)) \frac{|\tau - \tau'|}{\tau + \tau'} \\
&= (\tilde{\lambda})^2 \int_{2\tau_r}^\infty dt \int_{2\tau_r-t}^{t-2\tau_r} d\delta \exp(-\tilde{\lambda}(t - 2\tau_r)) \frac{|\delta|}{t} \\
&= 2(\tilde{\lambda})^2 \int_{2\tau_r}^\infty dt \int_0^{t-2\tau_r} d\delta \exp(-\tilde{\lambda}(t - 2\tau_r)) \frac{\delta}{t} \\
&= (\tilde{\lambda})^2 \int_{2\tau_r}^\infty dt \exp(-\tilde{\lambda}(t - 2\tau_r)) \frac{(t - 2\tau_r)^2}{t}
\end{aligned} \tag{A.9}$$

To analyze the limit of the C_{V2} for small refractory periods τ_r , we apply a Taylor expansion up to linear order.

$$\begin{aligned}
C_{V2}|_{\tau_r=0} &= \lambda \int_0^\infty dt \lambda \exp(-\lambda t) t = 1 \\
\frac{d}{d\tau_r} C_{V2} \Big|_{\tau_r=0} &= \lambda \int_0^\infty dt \lambda \exp(-\lambda t) (-\lambda^2 t^2 + 4\lambda t - 4) = -2\lambda \\
&\Rightarrow C_{V2} = 1 - 2\lambda\tau_r + \mathcal{O}((\tau_r)^2)
\end{aligned} \tag{A.10}$$

A.5 Clipped firing rates

Here, we summarize the derivations of the clipped firing rates omitting the trivial calculations for the Poisson Process (eq. 2.75) and for the PPR if $b \leq \tau_r$. For the PPR holds if $b \geq \tau_r$:

$$\begin{aligned}
\lambda_{\text{PPR}}^{\text{clip}} &= \frac{1}{b} \int_0^b dt \rho_1(t) \\
&= \frac{1}{b} \lambda \tau_r + \frac{1}{b} \int_{\tau_r}^b dt \lambda \exp(-\tilde{\lambda} \cdot (t - \tau_r)) \\
&= \frac{1}{b} \lambda \tau_r + \frac{1}{b} \frac{\lambda}{\tilde{\lambda}} (1 - \exp(-\tilde{\lambda}(b - \tau_r))) \\
&\stackrel{2.61}{=} \lambda \frac{\tau_r}{b} \exp(-\tilde{\lambda}(b - \tau_r)) + \frac{1}{b} (1 - \exp(-\tilde{\lambda}(b - \tau_r)))
\end{aligned} \tag{A.11}$$

For the gamma processes we get using the incomplete gamma functions $\Gamma(\gamma, x) = \int_x^\infty ds s^{\gamma-1} e^{-s}$ and $\bar{\Gamma}(\gamma, x) = \int_0^x ds s^{\gamma-1} e^{-s}$:

$$\begin{aligned}
\lambda_\gamma^{\text{clip}} &= \frac{1}{b} \int_0^b dt \rho_1^\gamma(t) \\
&= \frac{1}{b} \int_0^b dt \lambda \frac{\Gamma(\gamma, \gamma \lambda t)}{\Gamma(\gamma)} \\
&= \frac{1}{\gamma b} \frac{1}{\Gamma(\gamma)} \int_0^{\gamma \lambda b} ds \Gamma(\gamma, s) \\
&\stackrel{\text{p.I.}}{=} \frac{1}{\gamma b} \frac{1}{\Gamma(\gamma)} \left(\gamma \lambda b \cdot \Gamma(\gamma, \gamma \lambda b) + \int_0^{\gamma \lambda b} ds s^\gamma e^{-s} \right) \\
&= \lambda \frac{\Gamma(\gamma, \gamma \lambda b)}{\Gamma(\gamma)} + \frac{1}{b} \frac{\bar{\Gamma}(\gamma + 1, \gamma \lambda b)}{\Gamma(\gamma + 1)}
\end{aligned} \tag{A.13}$$

Appendix B

Calculations for Surrogate Methods

B.1 Effects of uniform dithering on the autocorrelation

The autocorrelation of an uniformly dithered surrogate spike train (with $\langle \rangle_{\delta_i}$ denoting the average according to the increment added to spike i):

$$\psi_{\delta}(t, t') := \sum_{i=1}^{\infty} \sum_{j=1}^{\infty} \langle \langle \delta(t - (t_i + \delta_i)) \delta(t' - (t_j + \delta_j)) \rangle_s \rangle_{\delta_i} \rangle_{\delta_j}$$

It is useful to split this sum up such that the average when $i = j$ is separated.

$$\begin{aligned} \psi_{\delta}(t, t') &= \sum_{i=1}^{\infty} \langle \langle \delta(t - (t_i + \delta_i)) \delta(t' - (t_i + \delta_i)) \rangle_s \rangle_{\delta_i} \\ &\quad + \sum_{i=1}^{\infty} \sum_{j=1, j \neq i}^{\infty} \langle \langle \delta(t - (t_i + \delta_i)) \delta(t' - (t_j + \delta_j)) \rangle_s \rangle_{\delta_i} \rangle_{\delta_j} \end{aligned}$$

We show that the first summand only depends on the firing rate $\lambda(t)$.

$$\begin{aligned}
& \sum_{i=1}^{\infty} \langle \langle \delta(t - (t_i + \delta_i)) \delta(t' - (t_i + \delta_i)) \rangle_s \rangle_{\delta_i} \\
&= \sum_{i=1}^{\infty} \langle \langle \delta(t - (t_i + \delta_i)) \rangle_s \rangle_{\delta_i} \delta(t - t') \\
&\stackrel{(2.4)}{=} \langle \langle s(t - \delta) \rangle_s \rangle_{\delta} \delta(t - t') \\
&\stackrel{(2.5)}{=} \langle \lambda(t - \delta) \rangle_{\delta} \delta(t - t') \\
&\stackrel{(5.2)}{=} \lambda_{\delta}(t) \delta(t - t')
\end{aligned}$$

While the second summand involves the autocorrelation function of the original process $\psi(t, t')$.

$$\begin{aligned}
& \sum_{i=1}^{\infty} \sum_{j=1, j \neq i}^{\infty} \langle \langle \langle \delta(t - (t_i + \delta_i)) \delta(t' - (t_j + \delta_j)) \rangle_s \rangle_{\delta_i} \rangle_{\delta_j} \\
&= \sum_{i=1}^{\infty} \sum_{j=1}^{\infty} \langle \langle \langle \delta(t - (t_i + \delta)) \delta(t' - (t_j + \delta')) \rangle_s \rangle_{\delta} \rangle_{\delta'} \\
&\quad - \sum_{i=1}^{\infty} \langle \langle \langle \delta(t - (t_i + \delta)) \delta(t' - (t_i + \delta')) \rangle_s \rangle_{\delta} \rangle_{\delta'} \\
&\stackrel{(2.4)}{=} \langle \langle \langle s(t - \delta) \rangle_s \rangle_{\delta} \rangle_{\delta'} \\
&\quad - \int ds \langle \langle \langle \delta(s - t_i) \delta(t - (s + \delta)) \delta(t' - (s + \delta')) \rangle_s \rangle_{\delta} \rangle_{\delta'} \\
&\stackrel{(2.7, 2.5)}{=} \langle \langle \psi(t - \delta, t' - \delta') \rangle_{\delta} \rangle_{\delta'} \\
&\quad - \int ds \langle \langle \lambda(s) \delta(t - (s + \delta)) \delta(t' - (s + \delta')) \rangle_{\delta} \rangle_{\delta'} \\
&\stackrel{(5.1)}{=} \int d\delta \int d\delta' u(\delta) u(\delta') \psi(t - \delta, t' - \delta') \\
&\quad - \int ds \int d\delta \int d\delta' \lambda(s) u(\delta) u(\delta') \delta(t - (s + \delta)) \delta(t' - (s + \delta')) \\
&= \int ds \int ds' u(t - s) u(t' - s') \underbrace{(\psi(s, s') - \lambda(s) \delta(s - s'))}_{:= \phi(s, s')}
\end{aligned}$$

$$\psi_{\delta}(t, t') = \lambda_{\delta}(t) \delta(t - t') + \int ds \int ds' u(t - s) u(t' - s') \phi(s, s')$$

Thus, it is possible to calculate $\psi_{\delta}(t)$ by using $\lambda(t)$ and $\psi(t, t')$.

B.2 Effects of uniform dithering on renewal processes

As we have seen in section 5.1, the convolution of the conditional firing rate with the p.d.f. of the uniform dithering become influential.

In the following we make extensively use of integral of the incomplete gamma function:

$$\int_a^b dx x^{n-1} \bar{\Gamma}(\gamma, x) \stackrel{\text{p.I.}}{=} \frac{1}{n} (x^n \bar{\Gamma}(\gamma, x) - \bar{\Gamma}(\gamma + n, x)) \Big|_{x=a}^{x=b}.$$

For gamma spike trains this is done within the following steps:

$$\begin{aligned} & (\lambda_c^+ \otimes u)(\tau) \\ &= \sum_{n=1}^{\infty} \frac{1}{2\Delta} \int_{-\Delta}^{\Delta} d\delta \gamma \lambda \frac{(\gamma \lambda (\tau - \delta))^{n\gamma-1}}{\Gamma(n\gamma)} \exp(-\gamma \lambda (\tau - \delta)) \Theta(\tau - \delta) \\ &= \frac{1}{2\Delta} \sum_{n=1}^{\infty} \int_{\max(0, \gamma \lambda (\tau - \Delta))}^{\max(0, \gamma \lambda (\tau + \Delta))} d\sigma \frac{\sigma^{n\gamma-1}}{\Gamma(n\gamma)} \exp(-\sigma) \\ &= \frac{1}{2\Delta} \sum_{n=1}^{\infty} \frac{1}{\Gamma(n\gamma)} (\bar{\Gamma}(n\gamma, \gamma \lambda (\tau + \Delta)) \Theta(\tau + \Delta) \\ &\quad - \bar{\Gamma}(n\gamma, \gamma \lambda (\tau - \Delta)) \Theta(\tau - \Delta)) \end{aligned}$$

$$(\lambda_c \otimes u)(0) = \frac{1}{2\Delta} \sum_{n=1}^{\infty} \frac{\bar{\Gamma}(n\gamma, \gamma \lambda \Delta)}{\Gamma(n\gamma)} = \frac{1}{2} (\lambda_c \otimes u)(0)$$

$$\begin{aligned}
& (\lambda_c^+ \circledast u \circledast u)(\tau) \\
&= \frac{1}{(2\Delta)^2} \sum_{n=1}^{\infty} \frac{1}{\Gamma(n\gamma)} \int_{-\Delta}^{\Delta} d\delta \left(\bar{\Gamma}(n\gamma, \gamma\lambda(\tau - \delta + \Delta)) \Theta(\tau - \delta + \Delta) \right. \\
&\quad \left. - \bar{\Gamma}(n\gamma, \gamma\lambda(\tau - \delta - \Delta)) \Theta(\tau - \delta - \Delta) \right) \\
&= \frac{1}{(2\Delta)^2} \sum_{n=1}^{\infty} \frac{1}{\Gamma(n\gamma)} \left[\left((\tau + 2\Delta) \bar{\Gamma}(n\gamma, \gamma\lambda(\tau + 2\Delta)) \right. \right. \\
&\quad \left. \left. - \frac{1}{\gamma\lambda} \bar{\Gamma}(n\gamma + 1, \gamma\lambda(\tau + 2\Delta)) \right) \Theta(\tau + 2\Delta) \right. \\
&\quad \left. - 2 \left(\tau \bar{\Gamma}(n\gamma, \gamma\lambda\tau) - \frac{1}{\gamma\lambda} \bar{\Gamma}(n\gamma + 1, \gamma\lambda\tau) \right) \Theta(\tau) \right. \\
&\quad \left. + \left((\tau - 2\Delta) \bar{\Gamma}(n\gamma, \gamma\lambda(\tau - 2\Delta)) \right. \right. \\
&\quad \left. \left. - \frac{1}{\gamma\lambda} \bar{\Gamma}(n\gamma + 1, \gamma\lambda(\tau - 2\Delta)) \right) \Theta(\tau - 2\Delta) \right]
\end{aligned}$$

$$\begin{aligned}
& (\lambda_c^+ \circledast u \circledast u)(0) \\
&= \frac{1}{(2\Delta)^2} \sum_{n=1}^{\infty} \frac{1}{\Gamma(n\gamma)} \left(2\Delta \bar{\Gamma}(n\gamma, 2\gamma\lambda\Delta) - \frac{1}{\gamma\lambda} \bar{\Gamma}(n\gamma + 1, 2\gamma\lambda\Delta) \right) \\
&= \frac{1}{(2\Delta)^2} \sum_{n=1}^{\infty} \left(2\Delta \frac{\bar{\Gamma}(n\gamma, 2\gamma\lambda\Delta)}{\Gamma(n\gamma)} - \frac{n}{\lambda} \frac{\bar{\Gamma}(n\gamma + 1, 2\gamma\lambda\Delta)}{\Gamma(n\gamma + 1)} \right) \\
&= \frac{1}{2} (\lambda_c \circledast u \circledast u)(0)
\end{aligned}$$

For the PPR spike trains the calculations are very similar:

$$\begin{aligned}
(\lambda_c^+ \otimes u)(\tau) &= \sum_{n=1}^{\infty} \frac{1}{2\Delta} \int_{-\Delta}^{\Delta} d\delta \left(\tilde{\lambda} \frac{(\tilde{\lambda}(\tau - \delta - n\tau_r))^{n-1}}{\Gamma(n)} \times \right. \\
&\quad \left. \times \exp(-\tilde{\lambda}(\tau - \delta - n\tau_r)) \Theta(\tau - \delta - n\tau_r) \right) \\
&= \frac{1}{2\Delta} \sum_{n=1}^{\infty} \int_{\max(0, \tilde{\lambda}(\tau - n\tau_r - \Delta))}^{\max(0, \tilde{\lambda}(\tau - n\tau_r + \Delta))} d\sigma \frac{\sigma^{n-1}}{\Gamma(n)} \exp(-\sigma) \\
&= \frac{1}{2\Delta} \sum_{n=1}^{\infty} \frac{1}{\Gamma(n)} \left(\bar{\Gamma}(n, \tilde{\lambda}(\tau - n\tau_r + \Delta)) \Theta(\tau - n\tau_r + \Delta) \right. \\
&\quad \left. - \bar{\Gamma}(n, \tilde{\lambda}(\tau - n\tau_r - \Delta)) \Theta(\tau - n\tau_r - \Delta) \right)
\end{aligned}$$

$$(\lambda_c^+ \otimes u)(0) = \frac{1}{2\Delta} \sum_{n=1}^{\lfloor \frac{\Delta}{\tau_r} \rfloor} \frac{1}{\Gamma(n)} \bar{\Gamma}(n, \tilde{\lambda}(\Delta - n\tau_r)) = \frac{1}{2} (\lambda_c \otimes u)(0)$$

$$\begin{aligned}
&(\lambda_c^+ \otimes u \otimes u)(\tau) \\
&= \frac{1}{(2\Delta)^2} \sum_{n=1}^{\infty} \frac{1}{\Gamma(n)} \times \\
&\quad \times \int_{-\Delta}^{\Delta} d\delta \left(\bar{\Gamma}(n, \tilde{\lambda}(\tau - n\tau_r - \delta + \Delta)) \Theta(\tau - n\tau_r - \delta + \Delta) \right. \\
&\quad \left. - \bar{\Gamma}(n, \tilde{\lambda}(\tau - n\tau_r - \delta - \Delta)) \Theta(\tau - n\tau_r - \delta - \Delta) \right) \\
&= \frac{1}{(2\Delta)^2} \sum_{n=1}^{\infty} \frac{1}{\Gamma(n)} \left[\left((\tau - n\tau_r + 2\Delta) \bar{\Gamma}(n, \tilde{\lambda}(\tau - n\tau_r + 2\Delta)) \right. \right. \\
&\quad \left. \left. - \frac{1}{\tilde{\lambda}} \bar{\Gamma}(n+1, \tilde{\lambda}(\tau - n\tau_r + 2\Delta)) \right) \Theta(\tau - n\tau_r + 2\Delta) \right. \\
&\quad \left. - 2 \left((\tau - n\tau_r) \bar{\Gamma}(n, \tilde{\lambda}(\tau - n\tau_r)) - \frac{1}{\tilde{\lambda}} \bar{\Gamma}(n+1, \tilde{\lambda}(\tau - n\tau_r)) \right) \Theta(\tau - n\tau_r) \right. \\
&\quad \left. + \left((\tau - n\tau_r - 2\Delta) \bar{\Gamma}(n, \tilde{\lambda}(\tau - n\tau_r - 2\Delta)) \right. \right. \\
&\quad \left. \left. - \frac{1}{\tilde{\lambda}} \bar{\Gamma}(n+1, \tilde{\lambda}(\tau - n\tau_r - 2\Delta)) \right) \Theta(\tau - n\tau_r - 2\Delta) \right]
\end{aligned}$$

$$\begin{aligned}
& (\lambda_c^+ \otimes u \otimes u)(0) \\
&= \frac{1}{(2\Delta)^2} \sum_{n=1}^{\lfloor \frac{2\Delta}{\tau_r} \rfloor} \frac{1}{\Gamma(n)} \left((2\Delta - n\tau_r) \bar{\Gamma}(n, \tilde{\lambda}(2\Delta - n\tau_r)) - \frac{1}{\tilde{\lambda}} \bar{\Gamma}(n+1, \tilde{\lambda}(2\Delta - n\tau_r)) \right) \\
&= \frac{1}{(2\Delta)^2} \sum_{n=1}^{\lfloor \frac{2\Delta}{\tau_r} \rfloor} \left((2\Delta - n\tau_r) \frac{\bar{\Gamma}(n, \tilde{\lambda}(2\Delta - n\tau_r))}{\Gamma(n)} - \frac{n}{\tilde{\lambda}} \frac{\bar{\Gamma}(n+1, \tilde{\lambda}(2\Delta - n\tau_r))}{\Gamma(n+1)} \right) \\
&= \frac{1}{2} (\lambda_c \otimes u \otimes u)(0)
\end{aligned}$$

B.3 Clipped firing rate for uniform dithering

How to calculate the clipped firing rate for a renewal process after uniform dithering? We know how the autocorrelation function looks like after uniform dithering (eq. (5.11)). Additionally, we have seen before that the autocorrelation function depends only on the firing rate and the conditional firing rate (eq. (A.2)) and that the conditional firing rate is linked to the ISI distribution (eq. (2.33)). This relation is invertible:

$$\lambda_c^+(\tau) = \sum_{i=1}^{\infty} p^{(i)\otimes}(\tau) \Leftrightarrow p(\tau) = \sum_{i=1}^{\infty} (-1)^{i-1} (\lambda_c^+)^{(i)\otimes}(\tau) \quad (\text{B.1})$$

Since $\lambda_c(\tau)$ is only defined for positive τ , it follows directly that $(\lambda_c^+)^{(i)\otimes}(\tau)$ is of order $\mathcal{O}(\tau^{i-1})$. Since we are only interested up to order $\mathcal{O}(\tau^2)$, we can write:

$$p(\tau) = \lambda_c^+(\tau) - (\lambda_c^+)^{(2)\otimes}(\tau) + \mathcal{O}(\tau^2) \text{ for } \tau > 0.. \quad (\text{B.2})$$

For this reasoning, it is necessary to assume renewalness of the process also after uniform dithering, that does not hold effectively, serial correlations appear, but the results for small τ have shown to fit very closely. How to get now the conditional firing rate from the autocorrelation after uniform dithering $\psi_\delta(t, t')$? We achieve this by

$$\lambda_c^{+, \delta}(\tau) = \frac{1}{\lambda} \psi_\delta(t, t + \tau) \text{ for } \tau > 0.$$

And thus:

$$\lambda_c^{+, \delta}(\tau) = (\lambda_c \otimes u \otimes u)(\tau) \text{ for } \tau > 0.$$

Expanding this up to order $\mathcal{O}(\tau^2)$ yields:

$$\lambda_c^{+, \delta}(\tau) = \underbrace{(\lambda_c \otimes u \otimes u)(0)}_{:= \lambda_u} + \mathcal{O}(\tau^2) \text{ for } \tau > 0.$$

This is due to the symmetry of $\bar{\lambda}_c$. Inserting this into (B.2) yields finally:

$$p_\delta(\tau) = (\lambda_u - (\lambda_u)^2 \tau + \mathcal{O}(\tau^2)) \Theta(\tau). \quad (\text{B.3})$$

B.4 Square root on the Joint-ISI distribution

In the papers by Gerstein (2004) and Louis et al. (2010), a square root is applied to the joint-ISI distribution. The following lines show why these papers use it and when avoiding it is possible.

The joint-ISI dithering is based on the joint-ISI distribution $f(\tau, \tau')$ (see sec. 5.3). For simplicity, we do not use the discretized time here; for small bin sizes b the differences of the results are negligible.

If we now want to dither the spike i , with its forward/backward ISIs τ_i/τ_{i+1} , the new position will be drawn according to (see eq. 5.19)

$$u(\delta) = \frac{\Theta(\Delta - |\delta|) f(\tau_i + \delta, \tau_{i+1} - \delta)}{\int_{-\Delta}^{\Delta} d\delta' f(\tau_i + \delta', \tau_{i+1} - \delta')}.$$

where $\theta(x)$ denotes the Heaviside step function, and Δ is the maximal displacement of the dithered spike.

Model calculation

To get a better understanding of the resulting joint-ISI histogram, we analyze here only the antidiagonal lines $D = ISI_i + ISI_{i+1} = \text{const}$. We assume the underlying distribution of the joint-ISI probability distribution is Gaussian on this antidiagonal, of mean zero and finite variance $\sigma^2 := \sigma^2(D)$. The assumption is to have a general ansatz that allows more analytical calculations than assuming PPR or gamma spike trains. The probability on the antidiagonal can thus be written as:

$$p_0(t) = c f(D/2 + t, D/2 - t) = \mathcal{N}(t, 0, \sigma^2). \quad (\text{B.4})$$

Here c is a normalization constant. Further, we use the notation:

$$\mathcal{N}(x, \mu, \sigma^2) := \frac{\exp(-\frac{(x-\mu)^2}{2\sigma^2})}{\sqrt{2\pi\sigma^2}}.$$

The distribution after the dithering process would then be $p_1(s) = \int p(s|t) p_0(t) dt$ with

$$p(s|t) = \frac{\Theta(\Delta - |s - t|) p_0(s)}{\int \Theta(\Delta - |s - t|) p_0(s) ds}.$$

An exact calculation for the transition probability yields:

$$p(s|t) = 2 \frac{\Theta(\Delta - |s - t|) \mathcal{N}(s, 0, \sigma^2)}{\operatorname{erf}(\frac{t+\Delta}{\sqrt{2}\sigma}) - \operatorname{erf}(\frac{t-\Delta}{\sqrt{2}\sigma})}.$$

Since an exact analytical calculation of $p_1(s)$ does not seem to be feasible, we will consider only the first two conditional moments of $p(s|t)$:

$$\langle s \rangle(t) = \int s \cdot p(s|t) ds = -2\sigma^2 \frac{\mathcal{N}(t, \Delta, \sigma^2) - \mathcal{N}(t, -\Delta, \sigma^2)}{\operatorname{erf}(\frac{t+\Delta}{\sqrt{2}\sigma}) - \operatorname{erf}(\frac{t-\Delta}{\sqrt{2}\sigma})},$$

$$\begin{aligned} \langle s^2 \rangle(t) &= \int s^2 \cdot p(s|t) ds \\ &= -2\sigma^2 \left(t \frac{\mathcal{N}(t, \Delta, \sigma^2) - \mathcal{N}(t, -\Delta, \sigma^2)}{\operatorname{erf}(\frac{t+\Delta}{\sqrt{2}\sigma}) - \operatorname{erf}(\frac{t-\Delta}{\sqrt{2}\sigma})} \right. \\ &\quad \left. + \Delta \frac{\mathcal{N}(t, \Delta, \sigma^2) + \mathcal{N}(t, -\Delta, \sigma^2)}{\operatorname{erf}(\frac{t+\Delta}{\sqrt{2}\sigma}) - \operatorname{erf}(\frac{t-\Delta}{\sqrt{2}\sigma})} \right) \\ &= t \cdot \langle s \rangle(t) - 2\sigma^2 \Delta \frac{\mathcal{N}(t, \Delta, \sigma^2) + \mathcal{N}(t, -\Delta, \sigma^2)}{\operatorname{erf}(\frac{t+\Delta}{\sqrt{2}\sigma}) - \operatorname{erf}(\frac{t-\Delta}{\sqrt{2}\sigma})}. \end{aligned}$$

Averaging over the initial positions yields

$$\langle s \rangle = \int \langle s \rangle(t) p_0(t) dt = 0,$$

since $\langle s \rangle(-t) = -\langle s \rangle(t)$ and $p_0(-t) = p_0(t)$. On the other hand, a numerical approximation is necessary for $\langle s^2 \rangle$.

Simplification of Model calculation

Since $\langle s^2 \rangle$ is not analytically calculable, we make a rough approximation here to achieve a comprehensible value for it. Thus, the window around the spike at time t to be dithered to time s : $h_\Delta(s - t) = \frac{1}{2\Delta} \Theta(\Delta - |s - t|)$, shall be changed from a rectangular to a normal distribution. In other words, h_Δ is approximated by its first two cumulants: $\tilde{h}_\Delta(s - t) = \mathcal{N}(s - t, 0, \sigma_\Delta^2)$ with $\sigma_\Delta^2 := \frac{\Delta^2}{3}$.

The corresponding transition probability is then:

$$\tilde{p}(s|t) = \frac{\tilde{h}_\Delta(s - t) p_0(s)}{\int \tilde{h}_\Delta(s - t) p_0(s) ds}$$

with

$$\tilde{h}_\Delta(s - t) p_0(s) = \mathcal{N}(s - t, 0, \sigma_\Delta^2) \cdot \mathcal{N}(s, 0, \sigma^2),$$

which yields:

$$\tilde{p}(s|t) = \mathcal{N}\left(s, \frac{\sigma^2}{\sigma^2 + \sigma_\Delta^2}t, \frac{\sigma^2\sigma_\Delta^2}{\sigma^2 + \sigma_\Delta^2}\right).$$

The resulting distribution is

$$\tilde{p}_1(s) = \mathcal{N}\left(s, 0, \sigma^2 \cdot \frac{\sigma^4 + \sigma^2\sigma_\Delta^2 + \sigma_\Delta^4}{(\sigma^2 + \sigma_\Delta^2)^2}\right). \quad (\text{B.5})$$

That means that after applying the joint-ISI dithering in this approximated case, we would still have a Gaussian distribution. In the two boundaries, we get

$$\lim_{\Delta \rightarrow 0} \tilde{p}_1(s) = \lim_{\Delta \rightarrow \infty} \tilde{p}_1(s) = p_0(s).$$

In this approximation, it is achieved thus that for both small Δ and large Δ , the resulting probability density corresponds to the original one - fulfilling the objective.

Effects of the square root

For intermediate window lengths Δ the variance of the distribution $\tilde{p}_1(s)$, i.e. the distribution after one entire joint-ISI dithering run, is unequal to the one of the original distribution $\text{Var}(t)_{\tilde{p}_1(t)} \neq \text{Var}(t)_{p_0(t)}$, which is a consequence of equations (B.5) & (B.5). Explicitly we have:

$$\text{Var}(t)_{\tilde{p}_1(t)} = \sigma^2 \cdot \frac{\sigma^4 + \sigma^2\sigma_\Delta^2 + \sigma_\Delta^4}{(\sigma^2 + \sigma_\Delta^2)^2} \neq \sigma^2 = \text{Var}(t)_{p_0(t)}$$

In the paper of 2004, George Gerstein proposed to apply a square root on the determined joint-ISI histogram. In the paper of Louis et al. 2010, it was argued the reason for this is a change in the kurtosis of the distribution, i.e., the fourth standardized moment, but they did not mention that also the variance changes, as we have seen above.

To overcome the change in the kurtosis, they applied a square root onto the joint-ISI distribution yielding the transition probability:

$$\begin{aligned} p_{\text{sqr}}(s|t) &= \frac{\tilde{h}_\Delta(s-t)\sqrt{p_0(s)}}{\int \tilde{h}_\Delta(s-t)\sqrt{p_0(s)}ds} \\ &= \mathcal{N}\left(s, \frac{\sigma^2}{\sigma^2 + \frac{1}{2}\sigma_\Delta^2}t, \frac{\sigma^2\sigma_\Delta^2}{\sigma^2 + \frac{1}{2}\sigma_\Delta^2}\right). \end{aligned}$$

The distribution after the dithering would follow as:

$$\begin{aligned} p_{1,\text{sqr}}(s) &= \int p_{\text{sqr}}(s|t)p_0(t)dt. \\ &= \mathcal{N}\left(s, 0, \sigma^2 \cdot \frac{\sigma^4 + \sigma^2\sigma_\Delta^2 + \frac{1}{2}\sigma_\Delta^4}{(\sigma^2 + \frac{1}{2}\sigma_\Delta^2)^2}\right). \end{aligned}$$

We know just compare the variances of the distributions, that are:

$$\begin{aligned}\text{Var}(t)_{\tilde{p}_1(t)} &:= \sigma_1^2 = \sigma^2 \cdot \frac{\sigma^4 + \sigma^2 \sigma_\Delta^2 + \sigma_\Delta^4}{(\sigma^2 + \sigma_\Delta^2)^2}, \\ \text{Var}(t)_{p_1, \tilde{\text{sqr}}(t)} &:= \sigma_2^2 = \sigma^2 \cdot \frac{\sigma^4 + \sigma^2 \sigma_\Delta^2 + \frac{1}{2} \sigma_\Delta^4}{(\sigma^2 + \frac{1}{2} \sigma_\Delta^2)^2}.\end{aligned}$$

Comparing them to the original variance σ^2 yields the result that the first distribution is narrower than the original one, and the second is broader than the original one, following:

$$\begin{aligned}\sigma_1^2 &= \sigma^2 \cdot \left(1 - \frac{\sigma^2 \sigma_\Delta^2}{(\sigma^2 + \sigma_\Delta^2)^2} \right) < \sigma^2, \\ \sigma_2^2 &= \sigma^2 \cdot \left(1 + \frac{\frac{1}{4} \sigma_\Delta^4}{(\sigma^2 + \frac{1}{2} \sigma_\Delta^2)^2} \right) > \sigma^2.\end{aligned}$$

For the boundary $\Delta \rightarrow 0$, σ_1^2 and σ_2^2 both tend against σ^2 . However, in this boundary, the spikes would not be displaced. It is thus not a useful application of a dithering method. In the other boundary $\Delta \rightarrow \infty$, σ_1^2 goes to σ^2 , while σ_2^2 goes to $2\sigma^2$. Thus, in the limit of long window lengths, if a square root is applied, the distribution becomes broader, while without this, the distribution after the dithering corresponds to the underlying original one. To reach this boundary Δ needs to be chosen such that $\sigma_\Delta^2 = \frac{\Delta^2}{3} \gg \sigma^2$. Since the broadness of the joint-ISI distribution is, in general depending, on the firing rate λ , in such a way that $\sigma^2 \propto \frac{1}{\lambda^2}$, it is difficult to choose a window length Δ such that for all firing rates the boundary situation is reached.

Summarizing, we understand now the reason for the change in the variance. Thus, we increase the window length used in Gerstein (2004) of $\Delta = 8$ ms to a default value in our implementation of $\Delta = 15$ ms (as for the other methods), which reduces the effect of the resulting distribution getting broader than the initial one.

Appendix C

Miscellaneous to False Positives and False Negatives

C.1 Pattern size distribution




Figure C.1: The figure shows the distribution of pattern sizes in the false positives analyzed in sec. 6.1. On the left and the right side, we show it for PPR spike trains ($\tau_r = 2.5$ ms) and gamma ($\gamma = 2$) spike trains. The distribution for the ground truth null hypothesis (GT) is shown in blue, using surrogates by uniform dithering (UD) in orange, by uniform dithering with refractoriness (UDR) in green, and by joint-ISI dithering (JISI) in red. The x -axis displays the pattern sizes and the y -axis the density.

We show here a further aspect of the analysis of independent data carried out in sec. 6.1. In this regard, we present the distribution of pattern sizes in the false positives (fig. C.1). We observe a clear tendency that false positive patterns contain a low number of spikes. Primarily, this holds for uniform dithering, where the distribution is built on a higher number of false positives. For this method, it is also worth mentioning that the size distribution varies only slightly between the underlying spike train types. Thus, we can deduce that the distribution of sizes in false positives may be independent of the spike train type.

C.2 P-value Spectra

In section 6.1, we have already shown the p-value spectra, that were a result of analyzing a set of PPR spike trains by recreating independent PPR spike trains, i.e., the ground truth, a by the use of uniform dithering (see fig. 6.3). Here, we show further the p-value spectra that arise using UDR and joint-ISI dithering in figure C.2. These two p-value spectra are very similar to the one of the ground truth null hypothesis and seem to be almost identical between each other. The fact that these two methods act very similarly applied to PPR spike trains also confirms this similarity (see sec. 5.4).

The p-value spectra that correspond to independent Gamma spike trains are shown in the figures C.3 & C.4. Comparing the p-value spectra of the ground truth between PPR (see fig. 6.3) and Gamma spike trains shows that it can be expected that the by-chance patterns in PPR spike trains have a bit higher occurrences. Regarding the surrogate methods, we observe that the uniform dithering also yields a shift to lower occurrences in the p-value spectrum. The same effect holds but with a smaller shift for the uniform dithering with refractoriness - only the p-value spectrum of the joint-ISI dithering fits closely to the ground truth one.

Summarizing, we can state that the p-value spectra in the case of PPR spike trains are well estimated using UDR or joint-ISI dithering, but in the case of Gamma spike trains only using joint-ISI dithering.

C.3 Transition from Poisson process to PPR

Here, we show that a Poisson process with refractoriness (PPR) results from a Poisson process deleting all spikes that have an interval to its preceding spike smaller than the refractory period τ_r . Given the Poisson spike train has the firing rate $\tilde{\lambda}$, the ISI distribution is:

$$p(\tau) = \tilde{\lambda} \exp(-\tilde{\lambda}\tau) \Theta(\tau).$$

Using the fact that the Poisson process is memoryless (see sec. 2.4.3.1), the ISI-distribution of the process that results from deleting the spikes which are too close to its predecessors yield the ISI-distribution:

$$p_r(\tau) \propto p(\tau)\Theta(\tau > \tau_r) = \tilde{\lambda} \exp(-\tilde{\lambda}\tau)\Theta(\tau - \tau_r).$$

The missing normalization factor is $\exp(\tilde{\lambda}\tau_r)$. Thus, it results in:

$$p_r(\tau) = \tilde{\lambda} \exp(-\tilde{\lambda}(\tau - \tau_r))\Theta(\tau - \tau_r),$$

which is the known ISI-distribution of the PPR process (see sec. 2.4.3.2).




Figure C.2: The figure shows the p-value spectrum corresponding to 20 independently generated PPR ($\tau_r = 2.5$ ms) spike trains of 1 s that were dithered by uniform dithering with refractoriness (upper panel) or with joint-ISI dithering (lower panel). In both cases, it is an average over the 100 data sets used to analyze the false positives of this specific firing rate. The three dimensions of the p-value spectrum are shown as such that every size has one subpanel. On the x -axis are the occurrences and on the y -axis the durations in bins. We show the p-values according to the logarithmic colormap on the right side.

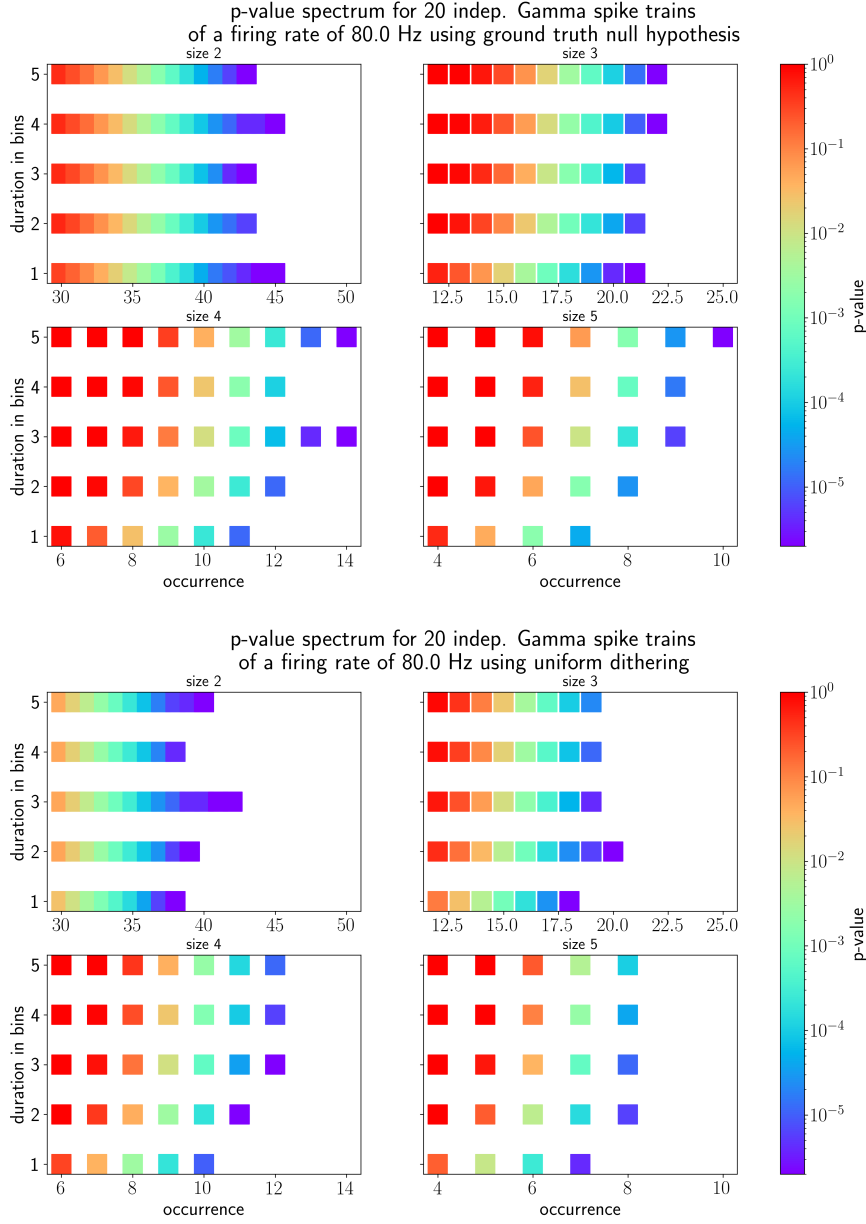


Figure C.3: The figure shows the p-value spectrum corresponding to 20 independently generated gamma ($\gamma = 2$) spike trains of 1 s (upper panel) and to uniformly dithered spike trains (lower panel). In both cases, it is an average over the 100 data sets used to analyze the false positives of this specific firing rate. The three dimensions of the p-value spectrum are shown as such that every size has one subpanel. On the x -axis are the occurrences and on the y -axis the durations in bins. We show the p-values according to the logarithmic colormap on the right side.

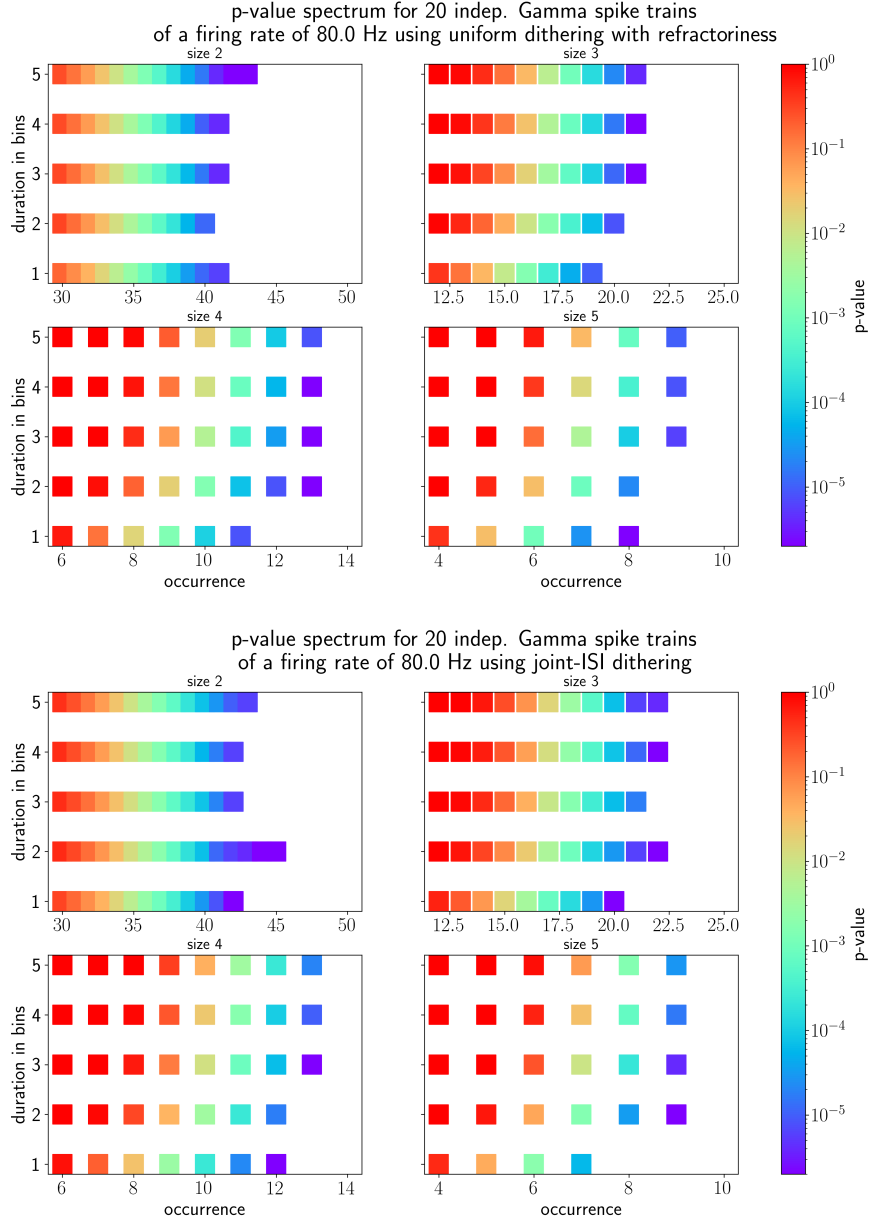


Figure C.4: The figure shows the p-value spectrum corresponding to 20 independently generated gamma ($\gamma = 2$) spike trains of 1 s that were dithered by uniform dithering with refractoriness (upper panel) or with joint-ISI dithering (lower panel). In both cases, it is an average over the 100 data sets used to analyze the false positives of this specific firing rate. The three dimensions of the p-value spectrum are shown as such that every size has one subpanel. On the x -axis are the occurrences and on the y -axis the durations in bins. We show the p-values according to the logarithmic colormap on the right side.

Bibliography

- Abeles, M. (1982). Role of the cortical neuron: integrator or coincidence detector? *Israel journal of medical sciences*, 18(1):83–92.
- Abeles, M. (1991). *Corticonics: Neural Circuits of the Cerebral Cortex*. Cambridge University Press, Cambridge, 1st edition.
- Abeles, M. (2010). Spatio-temporal patterns. In Rotter, S. and Grün, S., editors, *Analysis of Parallel Spike Trains*. Springer, Berlin.
- Asmussen, S. (2003). *Applied Probability and Queues*, volume 51 of *Stochastic Modelling and Applied Probability*. Springer, New York, 2 edition.
- 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. *J. Neurosci.*, 26(16):4166–4177.
- Benjamini, Y. and Hochberg, Y. (1995). Controlling the false discovery rate: A practical and powerful approach to multiple testing. *Journal of the Royal Statistical Society. Series B (Methodological)*, 57(1):289–300.
- Bienenstock, E. (1995). A model of neocortex. *Network: Computation in neural systems*, 6(2):179–224.
- 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.
- Braitenberg, V. and Schüz, A. (1991). *Anatomy of the Cortex: Statistics and Geometry*. Springer-Verlag, Berlin, Heidelberg, New York.
- Brette, R. (2015). Philosophy of the spike: Rate-based vs. spike-based theories of the brain. *Frontiers in Systems Neuroscience*, 9:151.
- Brochier, T., Zehl, L., Hao, Y., Duret, M., Sprenger, J., Denker, M., Grün, S., and Riehle, A. (2018). Massively parallel recordings in macaque motor cortex during an instructed delayed reach-to-grasp task. *Scientific Data*, 5:180055.

- Cardanobile, S. and Rotter, S. (2010). Simulation of stochastic point processes with defined properties. In Rotter, S. and Grün, S., editors, *Analysis of Parallel Spike Trains*. Springer, Berlin.
- Cox, D. R. and Isham, V. (1980). *Point Processes*. Number 12 in Monographs on Applied Probability and Statistics. Chapman & Hall, London.
- Cox, D. R. and Lewis, P. A. W. (1966). *The Statistical Analysis of Series of Events*. Methuen’s Monographs on Applied Probability and Statistics. Methuen, London.
- Cunningham, J. P., Gilja, V., Ryu, S. I., and Shenoy, K. V. (2009). Methods for estimating neural firing rates, and their application to brain-machine interfaces. *Neural Networks*, 22(9):1235–1246. Brain-Machine Interface.
- Daley, D. J. and Vere-Jones, D. (1988). *An Introduction to the Theory of Point Processes*. Springer Series in Statistics. Springer-Verlag, New York.
- Date, A., Bienenstock, E., and Geman, S. (1998). On the temporal resolution of neural activity. Technical report, Division of Applied Mathematics, Brown University.
- Dayan, P. and Abbott, L. F. (2001). *Theoretical Neuroscience*. MIT Press, Cambridge.
- Deger, M., Helias, M., Boucsein, C., and Rotter, S. (2012). Statistical properties of superimposed stationary spike trains. *Journal of Computational Neuroscience*, 32(3):443–463.
- Fino, E., Paille, V., Cui, Y., Morera-Herreras, T., Deniau, J.-M., and Venance, L. (2010). Distinct coincidence detectors govern the corticostriatal spike timing-dependent plasticity. *Journal of Physiology*, 588(16):3045–3062.
- Georgopoulos, A., Schwartz, A., and Kettner, R. (1986). Neuronal population coding of movement direction. *Science*, 4771(233):1416–1419.
- Gerstein, G. L. (2004). Searching for significance in spatio-temporal firing patterns. *Acta Neurobiol. Exp.*, 64:203–207.
- Gerstner, W., Kistler, W. M., Naud, R., and Paninski, L. (2014). *Neuronal Dynamics. From single Neurons to Networks and Models of Cognition*. Cambridge University Press, Cambridge.
- 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. nonstationary data. *Neural computation*, 14(1):81–119.
- Grün, S., Diesmann, M., and Aertsen, A. (2010). Unitary event analysis. In Rotter, S. and Grün, S., editors, *Analysis of Parallel Spike Trains*. Springer, Berlin.

- Harris, K. (2005). Opinion: Neural signatures of cell assembly organization. *Nature reviews. Neuroscience*, 6:399–407.
- Hatsopoulos, N., Geman, S., Amarasingham, A., and Bienenstock, E. (2003). At what time scale does the nervous system operate? *Neurocomputing*, (52–54):25–29.
- Hebb, D. O. (1949). *The organization of behavior: A neuropsychological theory*. John Wiley & Sons, New York.
- Heeger, D. (2000). Poisson model of spike generation.
- Helias, M., Tetzlaff, T., and Diesmann, M. (2015). Fluctuations and correlations in neuronal networks. Lecture Script, Jülich.
- Holm, S. (1979). A simple sequentially rejective multiple test procedure. *Scandinavian Journal of Statistics*, 6(2):65–70.
- Hong, S., Ratté, S., Prescott, S. A., and De Schutter, E. (2012). Single neuron firing properties impact correlation-based population coding. *Journal of Neuroscience*, 32(4):1413–1428.
- Izhikevich, E. M. (2006). Polychronization: computation with spikes. *Neural computation*, 18(2):245–282.
- Kandel, E. R., Schwartz, J. H., Jessel, T. M., Siegelbaum, S. A., and Hudspeth, A. J. (2015). *Principles of Neural Science*. McGraw-Hill, New York, 5 edition.
- Kass, R. E., Eden, U., and Brown, E. N. (2014). *Analysis of Neural Data*. Springer Series in Statistics. Springer.
- Kilavik, B. E., Roux, S., Ponce-Alvarez, A., Confais, J., Grün, S., and Riehle, A. (2009). Long-term modifications in motor cortical dynamics induced by intensive practice. *J. Neurosci.*, 29(40):12653–12663.
- Lamme, V. A., Super, H., Spekreijse, H., et al. (1998). Feedforward, horizontal, and feedback processing in the visual cortex. *Curr. Op. Neurobiol.*, 8(4):529–535.
- Louis, S., Borgelt, C., and Grün, S. (2010a). Generation and selection of surrogate methods for correlation analysis. In Rotter, S. and Grün, S., editors, *Analysis of Parallel Spike Trains*. Springer, Berlin.
- Louis, S., Gerstein, G. L., Grün, S., and Diesmann, M. (2010b). Surrogate spike train generation through dithering in operational time. *Front. Comput. Neurosci.*, 4(127).
- Nádasdy, Z., Hirase, H., Czurkó, A., Csicsvari, J., and Buzsáki, G. (1999). Replay and time compression of recurring spike sequences in the hippocampus. *J. Neurosci.*, 19(21):9497–9507.

- Nawrot, M. P., Boucsein, C., Rodriguez Molina, V., Riehle, A., Aertsen, A., and Rotter, S. (2008). Measurement of variability dynamics in cortical spike trains. *Journal of Neuroscience Methods*, 169:374–390.
- Onken, A., Grünwälder, S., Munk, M. H. J., and Obermayer, K. (2009). Analyzing short-term noise dependencies of spike-counts in macaque prefrontal cortex using copulas and the flashlight transformation. *PLOS Comput. Biol.*, 5(11):1–13.
- Pazienti, A. and Grün, S. (2006). Robustness of the significance of spike synchrony with respect to sorting errors. *Journal of Computational Neuroscience*, 21(3):329–342.
- Perez-Orive, J., Bazhenov, M., and Laurent, G. (2004). Intrinsic and circuit properties favor coincidence detection for decoding oscillatory input. *Journal of Neuroscience*, 24:6037–47.
- Perkel, D. H., Gerstein, G. L., and Moore, G. P. (1967). Neuronal spike trains and stochastic point processes. I. The single spike train. *Biophys. J.*, 7(4):391–418.
- Picado-Muiño, D., Borgelt, C., Berger, D., Gerstein, G. L., and Grün, S. (2013). Finding neural assemblies with frequent item set mining. *Front. Neuroinformatics*, 7:9.
- Pipa, G., Grün, S., and van Vreeswijk, C. (2013). Impact of spike train autostructure on probability distribution of joint spike events. *Neural Comput.*, 25.
- Prut, Y., Vaadia, E., Bergman, H., Haalman, I., Slovin, H., and Abeles, M. (1998). Spatiotemporal structure of cortical activity: properties and behavioral relevance. *J. Neurophysiol.*, 79(6):2857–2874.
- Quaglio, P. (2019). *Detection and Statistical Evaluation of Spike Patterns in Parallel Electrophysiological Recordings*. PhD thesis, RWTH Aachen University.
- Quaglio, P., Rostami, V., Torre, E., and Grün, S. (2018). Methods for identification of spike patterns in massively parallel spike trains. *Biol. Cybern.*, pages 1–24.
- Quaglio, P., Yegenoglu, A., Torre, E., Endres, D. M., and Grün, S. (2017). Detection and evaluation of spatio-temporal spike patterns in massively parallel spike train data with SPADE. *Frontiers in computational neuroscience*, 11:41.
- Rajdl, K. and Lansky, P. (2014). Fano factor estimation. *Mathematical biosciences and engineering : MBE*, 11:105–123.
- Rey, H. G., Pedreira, C., and Quiñero, R. (2015). Past, present and future of spike sorting techniques. *Brain Research Bulletin*, 119:106–117.
- Riehle, A., Brochier, T., Nawrot, M., and Grün, S. (2018). Behavioral context determines network state and variability dynamics in monkey motor cortex. *Front. Neural Circuits*, 12(52).

- 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.
- 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. *Nat. Rev. Neurosci.*, 7(9):982–991.
- Roy, S. A. and Alloway, K. D. (2001). Coincidence detection or temporal integration? what the neurons in somatosensory cortex are doing. *J. Neurosci.*, 21(7):2462–2473.
- Russo, E. and Durstewitz, D. (2017). Cell assemblies at multiple time scales with arbitrary lag constellations. *eLife*, 6:e19428.
- Salman, A., Malony, A., and Sottile, M. (2009). *An Open Domain-Extensible Environment for Simulation-Based Scientific Investigation (ODESSI)*. PhD thesis, University of Oregon.
- Shinomoto, S. (2010). Estimating the firing rate. In Rotter, S. and Grün, S., editors, *Analysis of Parallel Spike Trains*. Springer, Berlin.
- 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.
- Sridhar, S. (2020). Evaluation of automatic spike sorting algorithms: Reproducible workflows for comparison with ground truth and surrogate data. Master’s thesis, RWTH Aachen University.
- Stark, E. and Abeles, M. (2009). Unbiased estimation of precise temporal correlations between spike trains. *Journal of Neuroscience Methods*, 179(1):90–100.
- Stella, A. (2017). Comparison of statistical methods for spatio-temporal patterns detection in multivariate point processes: an application to neuroscience. Master’s thesis, University of Torino.
- Stella, A., Quaglio, P., Torre, E., and Grün, S. (2019). 3d-SPADE: Significance evaluation of spatio-temporal patterns of various temporal extents. *Biosystems*, 185:104022.
- 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. (2016). Synchronous spike patterns in macaque motor cortex during an instructed-delay reach-to-grasp task. *Journal of Neuroscience*, 36(32):8329–8340.
- van Vreeswijk, C. (2010). Stochastic models of spike trains. In Rotter, S. and Grün, S., editors, *Analysis of Parallel Spike Trains*. Springer, Berlin.

- Villa, A. E. and Abeles, M. (1990). Evidence for spatiotemporal firing patterns within the auditory thalamus of the cat. *Brain Res.*, 509(2):325–327.
- Yegenoglu, A., Quaglio, P., Torre, E., Grün, S., and Endres, D. (2016). Exploring the usefulness of formal concept analysis for robust detection of spatio-temporal spike patterns in massively parallel spike trains. In *International Conference on Conceptual Structures*, pages 3–16. Springer.
- Zaki, M. J. (2004). Mining non-redundant association rules. *Data Mining and Knowledge Discovery*, 9(3):223–248.

List of Figures

1.1	Schematic represensation of a neuron.	8
1.2	Temporal evolution of membrane potential; File:Action potential.svg. (2020, March 4). Wikimedia Commons, the free media repository. Retrieved 14:10, March 19, 2020 from https://commons.wikimedia.org/w/index.php?title=File:Action_potential.svg&oldid=401461195	8
1.3	Schematic representation of cell assemblies	10
1.4	Spatio-temporal patterns in a raster plot of spike trains. Figure adapted from Quaglio et al. (2018).	11
2.1	ISI distributions for PPR and gamma processes	25
2.2	Clipped firing rate for PPR and gamma processes	28
3.1	Firing rate profiles (reach-to-grasp)	31
3.2	Statistical Measures (reach-to-grasp)	32
4.1	Workflow of SPADE.	33
4.2	Scheme of an extracted spatio-temporal pattern	35
4.3	3-d p-value spectrum	37
5.1	Sketch explaining the uniform dithering method	40
5.2	Autocorrelation function for PPR and gamma processes before and after uniform dithering	43
5.3	Clipped firing rate for PPR and gamma processes (uniform dithering)	44
5.4	Sketch: Transition uniform dithering to uniform dithering with refractoriness	45
5.5	Scheme of the joint-ISI dithering	47
5.6	Effect of surrogate methods on firing rate profile	50
5.7	Cross-correlation function for PPR and gamma processes for different surrogate techniques	51
5.8	Rate of identical spikes for PPR and gamma processes for different surrogate techniques	52
5.9	Resulting ISI distributions for different surrogate methods	53
5.10	Autocorrelation function for PPR and gamma processes for different surrogate methods	55

5.11	Serial correlation coefficients for PPR and gamma processes after dithering	56
5.12	Clipped firing rate for PPR and gamma processes for different surrogate methods	57
5.13	Clipped firing rate for PPR and gamma processes for different surrogate methods	59
6.1	False Positive rate for ground truth null hypothesis	63
6.2	False Positive rate for different surrogate techniques	65
6.3	Comparison of ground truth and uniform dithering p-value spectra	67
6.4	Comparison of p-value spectra as contour plots	68
6.5	False Positive rate for ground truth null hypothesis	69
6.6	False Negative rates depending on the surrogate technique	73
6.7	False Positive rate for different surrogate techniques	75
6.8	False Discovery rate for different surrogate techniques	76
C.1	Distribution of pattern sizes in False positives	101
C.2	Comparison of UDR and joint-ISI p-value spectra - PPR spike trains . . .	104
C.3	Comparison of ground truth and uniform dithering p-value spectra	105
C.4	Comparison of UDR and joint-ISI p-value spectra - gamma spike trains .	106