



HAL
open science

Continuous testing for Poisson process intensities: a new perspective on scanning statistics

Franck Picard, Patricia Reynaud-Bouret, Etienne Roquain

► To cite this version:

Franck Picard, Patricia Reynaud-Bouret, Etienne Roquain. Continuous testing for Poisson process intensities: a new perspective on scanning statistics. *Biometrika*, 2018, 10.1093/biomet/asy044 . hal-02350783

HAL Id: hal-02350783

<https://hal.science/hal-02350783>

Submitted on 28 Mar 2024

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

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

Continuous testing for Poisson process intensities : A new perspective on scanning statistics

Franck Picard*

* *Univ Lyon, Université Lyon 1, CNRS, LBBE UMR5558,
F-69622 Villeurbanne, France
franck.picard@univ-lyon1.fr*

Patricia Reynaud-Bouret[‡]

[‡] *Université Côte d'Azur, CNRS, LJAD, France,
Patricia.Reynaud-Bouret@unice.fr*

Etienne Roquain[†]

[†] *LPMA, UMR 7599, Sorbonne Universités, UPMC, Univ. Paris 6,
F-75005, Paris, France.
etienne.roquain@upmc.fr*

May 25, 2017

Abstract

We propose a novel continuous testing framework to test the intensities of Poisson Processes. This framework allows a rigorous definition of the complete testing procedure, from an infinite number of hypothesis to joint error rates. Our work extends traditional procedures based on scanning windows, by controlling the family-wise error rate and the false discovery rate in a non-asymptotic manner and in a continuous way. The decision rule is based on a p -value process that can be estimated by a Monte-Carlo procedure. We also propose new test statistics based on kernels. Our method is applied in Neurosciences and Genomics through the standard test of homogeneity, and the two-sample test.

Contents

1	Introduction	4
2	The continuous testing framework for Poisson Processes	7
2.1	Homogeneity and two-sample conditional tests	7
2.2	An infinite set of local null hypothesis	8
2.3	A continuum of scanning windows	9
2.4	A first simple construction of the p -value process	9
2.5	Making continuous testing computationally tractable in general . .	10
2.6	General definition of the p -value process	12
3	False positive control in continuous time	12
3.1	Definition of the multiplicity error rates	13
3.2	A min- p procedure to control the FWER in continuous time	13
3.3	Link with the scan statistic framework.	14
3.4	A weighted BH procedure to control the FDR in continuous time .	15
4	Kernel-based statistics for sliding windows	17
4.1	On the limitations of the count-based statistic	17
4.2	The kernel vision in the two-sample case	18
5	Monte-carlo procedures for p-values estimation and adjustment	19
5.1	The distribution of the single test statistic being explicit.	19
5.2	The distribution of the single test statistic being not explicit	19
5.3	Estimation of the p -value process by Monte-Carlo	20
5.4	Two-step Monte-Carlo method for the min- p procedure.	21
6	Simulations	22
6.1	Homogeneity test	22
6.2	Two sample test	23
7	Applications	25
7.1	Do neurons respond to some stimulus ?	25
7.2	Differential analysis of replication origins between cell lines . . .	26
8	Conclusion	30
A	Supplementary material	35
A.1	Step-down improvement	35
A.2	Other kernel-based single test statistics	36
A.2.1	Homogeneity test statistics	36
A.2.2	Two-sample test statistics (one-sided)	37

B Proofs	38
B.1 Proof of Theorem 3.1	38
B.2 Proof of Theorem 3.5	38
B.3 Proof of Theorem 5.1	39

1 Introduction

Continuous testing has recently emerged as a suitable theoretical framework to test a possibly infinite number of hypothesis. It is especially adapted to situations where observed data come from an underlying continuously-indexed random process, like in the classical white noise model (Dümbgen and Spokoiny, 2001), or when data can be modeled by a point process (Blanchard et al., 2014). Our work being motivated by data coming from Neurosciences and Genomics, we focus on point processes that have been very effective to model occurrences of events through time or space in many fields. More than a broad range of applications, point processes have the advantage to be discrete by essence, which allows a drastic reduction of the continuum number of hypotheses, into to a random finite number of tests. This property was exploited by Blanchard et al. (2014) in a particular context, and we elaborate upon this work to propose a continuous testing framework that allows the rigorous definition of the complete testing procedure, from multiple hypothesis to error rates. Our procedure is very general and fully operational thanks to the use of sliding windows. This allows us to we revisit the old-tale of scanning statistics, and we extend their traditional possibilities and performance by providing the theoretical controls for the Family Wise Error Rate (FWER) as well as for the False Discovery Rate (FDR). We focus on Poisson processes for their remarkable properties, but the main definitions, methodological developments or computational tricks are valid for general point processes.

Scanning windows probably constitute the most common method to perform tests on Poisson processes for the detection of unexpected clusters of observations (Kulldorf, 1997; Chan and Zhang, 2007; Perone Pacifico et al., 2007; Siegmund et al., 2011) or to compare samples (Walther, 2010). It consists in computing a test statistic on overlapping windows of a given length that are shifted by all possible delays, which has the advantage of avoiding empirical discretization of the data (Grün, 1996; Tuleau-Malot et al., 2014). Then the decision rule is based on the so-called scan statistic, that refers to the window with most signal. Therefore these procedures rely on the distribution of the maximum of the test statistics under the full null, that can be either approximated (Naus, 1982; Robin et al., 2005; Fu et al., 2012; Rivera and Walther, 2013; Walther, 2010; Chan and Zhang, 2007) or learned by a randomization procedure (Kulldorff et al., 2005, 2009; Arias-Castro et al., 2015), which is common in standard multiple testing procedures (Westfall and Young, 1993; Romano and Wolf, 2005a). In this work we adopt the conditional testing approach (Loader, 1991; Rivera and Walther, 2013) that has the advantage of getting rid of nuisance parameters. It consists in deriving tests that are conditional to the *number* of occurrence (for the homogeneity test) or to the *position* of occurrences (for the two-sample test). Also, this strategy is consistent with the strong connections between continuous testing and non-parametric statistics (Blanchard et al., 2014), since the conditional test may depend on the unknown intensity functions of the observed processes. Therefore, our strategy is also to go beyond standard parametric approaches, that are mainly based on observed counts

or on a likelihood ratio statistic. As we will show, these parametric statistics do not fully exploit the spatial/temporal information that lie within the data, but only use integrated information. Inspired by recent results on non-parametric testing (Gretton et al., 2012; Fromont et al., 2013), we enrich the scan-statistics toolbox with kernel-based statistics to increase the power of detection of sharp alternatives.

Another limitation of current window-based strategies is that they correspond to global procedures that focus on a global null hypothesis. This resumes to ask whether the process is globally homogeneous, or if the samples are globally identical, which inevitably results in a global yes/no answer. However, with the increasing availability of more and more complex data, there is a need for flexible and local approaches to improve sensitivity for local structures. As we will see, continuous testing allows us to revisit the scanning windows framework, by defining a (possibly) uncountably infinite set of local hypothesis, and most importantly by providing a rigorous framework for defining joint error rates.

Indeed, the well known statistical challenge that arises when dealing with sliding windows is the appropriate control of Type-I errors, which is made difficult by the induced dependency structure. Current strategies mostly rely on the asymptotic approximation of the scan distribution (Kulldorf, 1997; Chan and Zhang, 2007), which has two advantages : *i*) it somewhat avoids the dependency matter by focusing on the scan statistic only, *ii*) it is computationally efficient. We rather adopt a non-asymptotic point of view, by considering the complete p -value process that becomes here a stochastic process . For some particular statistics (the count statistic for instance), this process is easily computed. When the statistic distribution is unknown we propose a Monte-Carlo procedure to estimate the p -value process. Then we propose a min- p procedure in continuous time to control the FWER in a non-asymptotic manner, based on a new computationally efficient double Monte-Carlo scheme. Simulations show that this non asymptotic procedure outperforms the one of Chan and Zhang (2007) in terms of FWER control.

The control of the False Discovery Rate (FDR), that is the expected proportion of errors among rejections, is more intricate, and tentatives are more recent. Perone Pacifico et al. (2004) propose a continuous analogue to the FDR in the context of random fields, where the amount of null hypotheses is assessed by the Lebesgue measure. However, their (continuous) FDR controlling procedure is based upon a FWER guarantee. Their strategy has also been considered in a setting close to our homogeneity framework (Perone Pacifico et al., 2007). Siegmund et al. (2011) also investigated the FDR for scanning statistics, but only under two strong conditions (Poisson distribution for the number of false discoveries and independence between null and alternative statistics), none of which being satisfied in standard models. Moreover, their procedure eventually resumes to discretized p -values , and is based on the tail approximation of the scan statistic, which is not valid for other statistics other than the scan statistic. Thanks to the continuous testing framework, we propose a weighted Benjamini-Hochberg procedure in continuous time for Poisson processes testing, which outperforms on simulations Siegmund et al. (2011). This procedure corresponds to a continuous analogue of the well-known

step-up procedure of Benjamini and Hochberg (1995), and extends the results of Blanchard et al. (2014) for continuous FDR in the context of local hypothesis testing.

The present work is inspired by two applications of Poisson processes: the analysis of spike trains data in Neurosciences and analysis of Genomic data coming from Next Generation Sequencing (NGS) technologies. In Neurosciences, Poisson processes are used to model spike trains, that is the succession of action potentials of a given neuron (Pipa et al., 2013). In particular inhomogeneous Poisson processes constitute a powerful model to understand rapid changes in firing rates induced by some stimuli (Kass et al., 2003). While the estimation of the Poisson process intensity has been already deeply studied (Shimazaki and Shinomoto, 2010), the questions asked in practice are often somehow different from a pure curve reconstruction matter: does the stimulus affect the spike apparition process (global hypothesis)? When (local) ? Does it depend on the type of stimulus and if so when exactly are the differences? In Genomics, next generation sequencing technologies (NGS) now allow the fine mapping of genomic features along chromosomes, and this spatial organization has become central to better understand genomes architecture and regulation. Here we consider the spatial organization of human replication origins that constitute the starting points of chromosomes duplication. To maintain the stability and integrity of genomes, replication origins are under a very strict spatio-temporal control, and part of their positioning has been shown to be associated with cell differentiation (Picard et al., 2014). Their recent mapping on the human genome has allowed the comparative analysis of replication origins maps between cell lines. Unfortunately, there is a lack of statistical framework to compare genomic maps of spatially organized genomic features, and we show how our procedure can be applied to the identification of replication starting sites that are linked to cell differentiation.

This article is organized as follows. First, we define the continuous testing framework, and we show how it can be easily handled in the Poisson process setup (Section 2), with a particular focus on the test of homogeneity and the two-sample test. We also define the infinite set of local test hypotheses, the continuum of scanning windows to test them and the associated p -value stochastic process. To proceed, and for the sake of simplicity, we first focus on the standard count statistic that directly gives access to the p -values. Then two types of control are discussed in Section 3: the classical FWER and the more involved control of the FDR in continuous time. Both procedures result in an adjusted p -value process, which makes our method operational and easily implemented. To increase the power of our procedure we propose in Section 4 local kernel-based statistics. Those kernel statistics require non explicit conditional distributions to be transformed into p -values and we discuss Monte-Carlo procedures to estimate the p -values, with a new computationally efficient double Monte-Carlo step to control the FWER in Section 5. In Section 6 we show through simulations that the resulting procedures are more powerful than existing ones (in particular those based on unconditioned scanning statistics or on discretization). Section 7.2 shows that the procedure works

extremely well on experimental data, issued from Neurosciences or Genomics. Our method is available in the form of the `contest` R-package, and offers many perspectives in the field of continuous testing.

2 The continuous testing framework for Poisson Processes

In the sequel a point process is a random countable set of points usually denoted by N . The corresponding point measure is then denoted by dN , that is $\sum_{T \in N} \delta_T$ with δ denoting the Dirac mass. For any interval I , $N(I)$ is a random variable that corresponds to the number of points of N in I . We focus on processes in \mathbb{R} (which can be thought as the DNA strand for genomic applications and time line for neuroscience applications) and we consider processes in $[0,1]$ for the sake of simplicity.

2.1 Homogeneity and two-sample conditional tests

In order to propose a unified framework, we use the following notations: we observe a random set of points Y with a distribution of the form $P_{\theta, \lambda}$, where λ is a nuisance parameter, while θ represents the signal of interest. In both cases, the aim is to test whether θ is locally null.

Homogeneity test. We observe $Y = N$, a point process modeled by a heterogeneous Poisson Process on $[0, 1]$ with unknown intensity ν (non negative and in $L^2[0, 1]$) with respect to the Lebesgue measure. We would like to detect the regions that are particularly rich or poor in occurrences of the process, with respect to its mean/stationary behavior. To do so, we compare the intensity ν of the observed process to the mean behavior of the intensity

$$\lambda = \int_0^1 \nu(s) ds.$$

Then we rewrite the distribution of Y in terms of two different parameters λ and θ , with θ , defined by

$$\forall t \in [0, 1], \theta(t) = \frac{\nu(t)}{\int_0^1 \nu(s) ds} - 1.$$

The signal of interest is θ , whereas λ may be seen as a nuisance parameter because its value is unknown under the null hypothesis. The homogeneity null hypothesis is therefore reduced to " $\theta = 0$ ": it is a composite hypothesis since under the null hypothesis, the distribution can still vary according to λ . As a particular easier problem, one can also assume that λ is known. The problem is then a classical goodness-of-fit test and the null hypothesis is not composite anymore.

Two-sample test. We observe here $Y = (N_A, N_B)$, where N_A and N_B are two heterogeneous independent Poisson Processes on $[0, 1]$ with intensities ν_A and ν_B (non negative and in $L^2[0, 1]$) with respect to the Lebesgue measure. In the following, for any interval I , we denote $Y(I) = (N_A(I), N_B(I))$, and by a slight abuse of notation, $(Y \cap I) = (N_A \cap I, N_B \cap I)$. Our aim here is to detect the regions where $\nu_A \neq \nu_B$ (or, alternatively, $\nu_A > \nu_B$).

To re-parametrize the distribution of Y , one can provide a one-to-one correspondence between Y and the couple (N, ε) , where $N = N_A \cup N_B$ is the joint process and where $\varepsilon = (\varepsilon_T)_{T \in N}$ is a set of marks with values in $\{-1, 1\}$ such that $\varepsilon_T = 1$ if T actually belongs to N_A and $\varepsilon_T = -1$ if T actually belongs to N_B (Kingman, 1993). An important point here is to note that N is actually a Poisson process with intensity $\lambda = \nu_A + \nu_B$ and that conditionally to N , the ε_T 's are independent variables with distribution

$$\varepsilon_T | N \sim 2\mathcal{B}\left(\frac{\theta(T) + 1}{2}\right) - 1,$$

with

$$\forall t \in [0, 1], \theta(t) = \frac{\nu_A(t) - \nu_B(t)}{\nu_A(t) + \nu_B(t)}.$$

Once again, θ stands for the amount of signal in the data and λ is a nuisance parameter. The two-sample null hypothesis is then once again summarized by " $\theta = 0$ ". Since λ is an unknown function, this hypothesis is composite. This case is thus much more complex than the homogeneity case, and calls for a non-parametric procedure.

Conditional distributions. To avoid in both cases the use of specific unknown values of λ to derive the distribution under the null hypothesis, our strategy is based on conditional testing (Loader, 1991; Rivera and Walther, 2013). Indeed, for the homogeneity test, conditionally on the total number of points, $N([0, 1]) = n$, the observed process is actually an n *i.i.d.* sample of density $1 + \theta$ that does not depend on λ (Loader, 1991; Rivera and Walther, 2013). Similarly, in the two-sample case, given the positions of the joint process N , the remaining variability under the null hypothesis lies in the distribution of the marks $(\varepsilon_T)_{T \in N}$ that only depends on θ (and not on λ anymore). Thus we introduce the variable \mathfrak{N} , such that $\mathfrak{N} = N([0, 1])$ in the homogeneity case, and $\mathfrak{N} = N$ in the two-sample case and our testing procedure relies on the distribution of Y conditionally to \mathfrak{N} . Since the conditional distribution of Y given \mathfrak{N} does not depend on λ , it is denoted by P_θ in the sequel, whereas E_λ refers to the expectation w.r.t. \mathfrak{N} .

2.2 An infinite set of local null hypothesis

The continuous testing procedure developed here aims at determining where the signal θ is non zero. The classical strategy consists in testing the *full* null hypothesis, which corresponds to the single question "Is the signal θ null over $[0, 1]$?".

Therefore a unique hypothesis is tested and the test produces a binary answer. Here, we rather consider local null hypothesis such as,

$$\begin{aligned} & \forall t \in [0, 1], H_{0,t} : \{\theta(t) = 0\} \quad \text{against } H_{1,t} : \{\theta(t) > 0\}, \\ \text{or } & \forall t \in [0, 1], H_{0,t} : \{\theta(t) = 0\} \quad \text{against } H_{1,t} : \{\theta(t) \neq 0\}. \end{aligned}$$

In practice, one-sided tests of homogeneity are usually considered to detect regions that are too rich in occurrences, whereas two-sided tests of the two-sample problem seem more relevant when there is a complete equivalence between both processes N_A and N_B from a modeling point of view.

Since information at the single point level seems difficult to catch in general, we focus on finding intervals or unions of intervals on which the null hypothesis holds. We denote $\mathcal{H}_0\{I\}$, the null hypothesis corresponding to “ $\forall t \in I, \theta(t) = 0$ ”. As a particular case, $\mathcal{H}_0\{[0, 1]\}$ is the full null hypothesis and corresponds to “ $\theta = \theta_0$ ”, where θ_0 is the null function on $[0, 1]$.

2.3 A continuum of scanning windows

Our procedure is based on scanning windows with a given resolution $\eta \in]0, 1[$ that is fixed. We envision here all the possible *continuum* of windows with such given length, meaning that we are performing a whole *continuum* of tests, one for each window. To avoid any confusion in the sequel, x always denotes a *window center* whereas t always denotes a *point*, that is a possible value for the observations Y . Thus we have $x \in \mathcal{X}_\eta =]\eta/2, 1 - \eta/2[$ (while the set of possible values for t is $[0, 1]$). The scanning windows are denoted by $I_\eta(x) = (x - \eta/2, x + \eta/2)$ for any $x \in \mathcal{X}_\eta$. Therefore, a typical relationship between x and t is “ $t \in I_\eta(x)$ ”. In the sequel, all the proposed multiple testing procedures are actually based on single tests of the null hypothesis $\mathcal{H}_0\{I_\eta(x)\}$, for all the possible window centers $x \in \mathcal{X}_\eta$, which is different than the null hypotheses $H_{0,t}$ given by $\theta(t) = 0$ for t in $[0, 1]$. The continuous testing procedure aims at accepting the set of “true windows”, which is indexed by their centers:

$$J_0^\eta(\theta) = \{x \in \mathcal{X}_\eta : \forall t \in I_\eta(x), \theta(t) = 0\}.$$

$J_0^\eta(\theta)$ formally depends on θ , but is denoted by J_0^η in the sequel, for short.

2.4 A first simple construction of the p -value process

Homogeneity test. Consider as an individual test statistic $S_\eta(x) = N(I_\eta(x))$, that is the number of points in the window of center x . Then classically, a one-sided test based on this statistic rejects $\mathcal{H}_0\{I_\eta(x)\}$ if $S_\eta(x)$ is large enough. Under $\mathcal{H}_0\{I_\eta(x)\}$, it is easy to see that the conditional distribution of $S_\eta(x)$ given $\mathfrak{N} = n$ (or $N([0, 1]) = n$ here) is a Binomial variable with parameter n and $1/2$. Hence, if we denote by $F_{\mathcal{B}(n, \frac{1}{2})}(\cdot)$, the function such that $F_{\mathcal{B}(n, \frac{1}{2})}(z) = P(Z \geq z)$ with

$Z \sim \mathcal{B}(n, \frac{1}{2})$, then

$$\forall x \in \mathcal{X}_\eta, p_\eta(x) = F_{\mathcal{B}(N([0,1]), \frac{1}{2})}(S_\eta(x)),$$

can be seen as a p -value associated to the classical one-sided single test of $\mathcal{H}_0 \{I_\eta(x)\}$. Since $F_{\mathcal{B}(n, \frac{1}{2})}$ is explicit and since the number of points, as a function of x , can be very efficiently computed, one can have access to a p -value process $x \mapsto p_\eta(x)$, with a very low computational cost.

If λ is known, one does not need to use a conditional distribution. The unconditional distribution of $S_\eta(x)$ is a Poisson distribution with mean $\eta\lambda$, under $\mathcal{H}_0 \{I_\eta(x)\}$. If $F_{\mathcal{P}(\eta\lambda)}(z) = P(Z \geq z)$ with $Z \sim \mathcal{P}(\eta\lambda)$, then

$$\forall x \in \mathcal{X}_\eta, p_\eta(x) = F_{\mathcal{P}(\eta\lambda)}(S_\eta(x)) \quad (1)$$

also defines a p -value process. Note that this set-up is quite close to the one of Chan and Zhang (2007).

Two-sample test. Consider now as individual test statistic $S_\eta(x) = N_A(I_\eta(x))$. A classical one-sided test based on this statistic also rejects $\mathcal{H}_0 \{I_\eta(x)\}$, if $S_\eta(x)$ is large enough. Under $\mathcal{H}_0 \{I_\eta(x)\}$, it is easy to see that the conditional distribution of $S_\eta(x)$ given $\mathfrak{N} = N$ (the joint process here), is the same as the conditional distribution of $S_\eta(x)$ given $N(I_\eta(x)) = N_A(I_\eta(x)) + N_B(I_\eta(x))$. It is a Binomial variable with parameter $N(I_\eta(x))$ and $1/2$. The corresponding p -value associated to the classical one-sided single test of $\mathcal{H}_0 \{I_\eta(x)\}$ can then be defined by:

$$\forall x \in \mathcal{X}_\eta, p_\eta(x) = F_{\mathcal{B}(N(I_\eta(x)), \frac{1}{2})}(S_\eta(x)). \quad (2)$$

This process depends on N and x , and once again can be computed very easily.

2.5 Making continuous testing computationally tractable in general

Even if the method is based on a *continuum* of tests, the two previous examples are computationally tractable because they are based on $Y(I_\eta(x))$, the number of points within windows $I_\eta(x)$, which can be computed very efficiently and which is piece-wise constant as a function of x . In full generality, the number of points may not be informative enough, and to gain power, one can consider more general single test statistics (see Section 4) that only depend on the *composition* of each window, that is on the function

$$x \in \mathcal{X}_\eta \mapsto Y \cap I_\eta(x),$$

which records the random positions of points in Y that lie in $I_\eta(x)$. Since the composition is piece-wise constant with changes only each time a point leaves or enters the scanning window, the single test statistic that only depends on the composition is a piece-wise constant function on a random finite partition $\tau = (\tau_m)_{0 \leq m \leq M}$. Therefore, the *continuum* of tests can be reduced to a random finite number of tests, one for each segment defined by τ . An example of such a partition is provided in Fig. 1.




Figure 1: Construction of the p -value process for the two-sample test. Occurrences from A and B are merged to form the joint process N and Rademacher marks $(\varepsilon_T^0)_T$ are introduced as labels. Each occurrence $T \in N$ has a span η that is used to create the partition τ whose elements constitute the centers of the testing windows $I_\eta(\tau_m)$, $(\tau_m \in \tau)$. Such construction of the windows ensures that the composition of the observed point process is constant between two centers τ_m and τ_{m+1} in terms of number and repartition of points. The p -value process is then a càdlàg process with jumps defined by partition τ .

2.6 General definition of the p -value process

In the sequel, we assume that a test statistic $S_\eta(x)$ is given for each window $I_\eta(x)$, that only depends on the composition of the window. Because Poisson processes have nice properties of independence between disjoint sets, it is easy to show the following property, which is usually referred to as subset pivotality (Westfall and Young, 1993; Romano and Wolf, 2005a).

$$(\mathcal{P}) \quad \left| \begin{array}{l} \text{For all } \theta, \text{ the conditional distribution of } (S_\eta(x))_{x \in J_0^n} \text{ given } \mathfrak{N} \\ \text{is the same whether } Y \sim P_\theta \text{ or } Y \sim P_{\theta_0}. \end{array} \right.$$

This roughly says that the conditional distribution of $S_\eta(x)$ given \mathfrak{N} is not modified by the truth or the falsehood of the remaining hypotheses. In particular, the conditional distribution of $S_\eta(x)$ given \mathfrak{N} under $\mathcal{H}_0 \{I_\eta(x)\}$ is the same as the one under the full null $\mathcal{H}_0 \{[0, 1]\}$.

Note that in full generality the conditional distribution of $S_\eta(x)$ given \mathfrak{N} depends on \mathfrak{N} but also on the composition of the window centered on x . Therefore we define:

$$\forall x \in J_0^n, F_{\mathfrak{N},x}(s) = P_{\theta_0}(S_\eta(x) \geq s | \mathfrak{N}) = P_\theta(S_\eta(x) \geq s | \mathfrak{N}),$$

and the associated p -value process by

$$p_\eta(x) = F_{\mathfrak{N},x}(S_\eta(x)), \quad x \in \mathcal{X}_\eta. \quad (3)$$

It satisfies for all $x \in \mathcal{X}_\eta$, for all (θ, λ) ,

$$\text{if } \mathcal{H}_0 \{I_\eta(x)\} \text{ is true, } \forall \alpha \in [0, 1], P_{\theta,\lambda}(p_\eta(x) \leq \alpha) \leq \alpha. \quad (4)$$

Moreover since $p_\eta(x)$ is a deterministic function of $S_\eta(x)$ (conditionally to \mathfrak{N}), property (\mathcal{P}) also holds for $p_\eta(x)$, and not only for $S_\eta(x)$.

3 False positive control in continuous time

Except when specifically mentioned, the results presented here hold for any test statistic $S_\eta(x)$ that only depends on the composition of the window (as defined in Section 2.5). Their associated p -value process are as described previously.

A rejection set \mathcal{R}_η is a subset of \mathcal{X}_η , which is a function of the p -value process $(p_\eta(x))_{x \in \mathcal{X}_\eta}$. Classically, it corresponds to small values of this process and a typical example is given by

$$\mathcal{R}_\eta(u) := \{x \in \mathcal{X}_\eta : p_\eta(x) \leq u\},$$

for a given threshold u potentially depending on the data. Once the rejection set \mathcal{R}_η is given, the set of accepted windows is defined by its complement $\mathcal{A}_\eta := (\mathcal{R}_\eta)^c$.

3.1 Definition of the multiplicity error rates

If the procedure behaves properly, the set \mathcal{A}_η should be a good approximation of J_0^η , or similarly \mathcal{R}_η should be a good approximation of $(J_0^\eta)^c$. To evaluate the quality of a procedure \mathcal{R}_η , we focus on *false positive windows* that correspond to elements of $J_0^\eta \cap \mathcal{R}_\eta$. The size of $J_0^\eta \cap \mathcal{R}_\eta$ can be gauged in many ways. In this work, we focus on the most famous criteria, namely the Family-Wise Error Rate (FWER) and the False Discovery Rate (FDR).

From a family-wise point of view, one wants to avoid any false positive with a large probability, which corresponds to a control of the quantity

$$\text{FWER}_{\theta,\lambda}^\eta(\mathcal{R}_\eta) = P_{\theta,\lambda}(J_0^\eta \cap \mathcal{R}_\eta \neq \emptyset). \quad (5)$$

While a FWER control provides a strong assessment on the amount of false positives, this approach generally leads to conservative procedures. A more balanced view is to allow some false positives, in a pre-specified fraction of the total amount of positives. This can be achieved by controlling

$$\text{FDR}_{\theta,\lambda}^\eta(\mathcal{R}_\eta) = E_{\theta,\lambda} \left(\frac{\Lambda(J_0^\eta \cap \mathcal{R}_\eta)}{\Lambda(\mathcal{R}_\eta)} \right), \quad (6)$$

with $0/0 = 0$. Importantly, since the setting is continuous here, the “number” of (false) positives is quantified above by the Lebesgue measure, Λ , on $[0, 1]$.

3.2 A min- p procedure to control the FWER in continuous time

We would like to find a threshold $u^\alpha \in [0, 1]$ such that:

$$\forall(\theta, \lambda), \text{FWER}_{\theta,\lambda}^\eta(\mathcal{R}_\eta(u^\alpha)) \leq \alpha.$$

For all $u \in [0, 1]$,

$$\{J_0^\eta \cap \mathcal{R}_\eta(u) \neq \emptyset\} = \{\exists x \in J_0^\eta : p_\eta(x) \leq u\} = \left\{ \inf_{x \in J_0^\eta} \{p_\eta(x)\} \leq u \right\},$$

with $p_\eta(x)$ defined in (3). The set J_0^η being unknown, one can use some rough upper-bound. We have, thanks to subset pivotality (\mathcal{P}), that for any $u > 0$ (possibly depending on the conditioning variable \mathfrak{N}),

$$\begin{aligned} \text{FWER}_{\theta,\lambda}^\eta(\mathcal{R}_\eta(u)) &= E_\lambda \left[P_\theta \left(\inf_{x \in J_0^\eta} \{p_\eta(x)\} \leq u \mid \mathfrak{N} \right) \right] \\ &= E_\lambda \left[P_{\theta_0} \left(\inf_{x \in J_0^\eta} \{p_\eta(x)\} \leq u \mid \mathfrak{N} \right) \right]. \end{aligned}$$

Since $J_0^\eta \subset \mathcal{X}_\eta$, we obtain

$$\text{FWER}_{\theta,\lambda}^\eta(\mathcal{R}_\eta(u)) \leq E_\lambda \left[P_{\theta_0} \left(\inf_{x \in \mathcal{X}_\eta} \{p_\eta(x)\} \leq u \mid \mathfrak{N} \right) \right].$$

The distribution of $\inf_{x \in \mathcal{X}_\eta} \{p_\eta(x)\}$ under θ_0 and given \mathfrak{N} are in both cases (homogeneity and two sample cases) known/observable once the variable \mathfrak{N} is fixed and it is possible to calibrate an adequate threshold u^α . However, due to the classical subtleties of inversion of non continuous c.d.f., it is in fact more powerful to use *adjusted p-values* that are defined by

$$q_\eta(x) = F_{\mathfrak{N}}^m(p_\eta(x)), \quad x \in \mathcal{X}_\eta, \quad (7)$$

with $F_{\mathfrak{N}}^m$ being the conditional c.d.f. of $\inf_{x \in \mathcal{X}_\eta} \{p_\eta(x)\}$ under the full null hypothesis, that is $\theta = \theta_0$.

Theorem 3.1. *Let $\alpha \in (0, 1)$ and q_η given by (7). Then the procedure defined by*

$$\mathcal{R}_\eta^m := \{x \in \mathcal{X}_\eta : q_\eta(x) \leq \alpha\} \quad (8)$$

has a FWER controlled by α .

Note that \mathcal{R}_η^m , which is a continuous version of the classical min- p procedure, is coherent with the case where there is only one value in the p -value process because the resolution η is equal to 1. In this extreme case, \mathcal{X}_η is reduced to one point and the p -value is exactly the adjusted p -value, the control of $\text{FWER}(\mathcal{R}_\eta^m)$ being then an obvious consequence of Section (4).

Remark 3.2. *The previous procedure is usually referred to as a single-step procedure. We also propose refinements by using a step-down procedure in the Supplementary Material.*

Remark 3.3. *Unfortunately, even for the classical homogeneity or two-sample tests the c.d.f. $F_{\mathfrak{N}}^m$ is not explicit and needs to be estimated. This is the purpose of Section 5 in which we also derive meaningful theoretical results when this distribution is approximated by Monte-Carlo that are necessary to provide a completely operational method with strong theoretical guarantee.*

3.3 Link with the scan statistic framework.

In the case of the homogeneity test, the p -value process does not depend on x and because function $F_{\mathcal{B}(N([0,1]), \frac{1}{2})}$ is a non increasing function, we have

$$\inf_{x \in \mathcal{X}_\eta} p_\eta(x) = F_{\mathcal{B}(N([0,1]), \frac{1}{2})} \left(\sup_{x \in \mathcal{X}_\eta} S_\eta(x) \right). \quad (9)$$

Hence, rejecting for small values of the infimum of the p -value process (similar to standard minP procedures), or rejecting for large of the supremum of the test statistic (i.e. the scan statistic, often called the maxT procedure) is strictly equivalent (Dudoit et al., 2003; Daley and Vere-Jones, 2008). Compared with the work of Chan and Zhang (2007), we mainly added two ingredients: (i) the distribution

we are considering is conditional and exact, whereas theirs is asymptotic and still depends on the unknown nuisance parameter λ (even if they provide a plug-in procedure in practice to avoid this issue), (ii) thanks to Property (\mathcal{P}), we are able to prove the FWER control, whereas their control stands under the full null only.

This minP-maxT equivalence is valid in the case of the homogeneity test only, but it does not hold as soon as the c.d.f. of the infimum of the p -value under the full null depends on the composition of the window. This is typically the case for the two-sample test, where $F_{\mathcal{B}(N(I_\eta(x)), \frac{1}{2})}$ depends on x so that Eq. (9) cannot hold anymore. In this case, the minP framework offers the advantage to scale all tests on the same level (Dudoit et al., 2003; Daley and Vere-Jones, 2008).

3.4 A weighted BH procedure to control the FDR in continuous time

To increase the detection capability, we propose a second procedure based on the control of the FDR criterion (6).

Heuristic. We would like to find a threshold $v^\alpha \in [0, 1]$ (see (3)) such that:

$$\forall(\theta, \lambda), \text{FDR}_{\theta, \lambda}^\eta(v^\alpha) \leq \alpha.$$

Following Blanchard et al. (2014), let us first explain heuristically how one can reach this goal in a continuous framework. By Fubini's theorem, $\forall v \in [0, 1]$,

$$\text{FDR}_{\theta, \lambda}^\eta(v) = E_{\theta, \lambda} \left(\frac{\Lambda(J_0^\eta \cap \mathcal{R}_\eta(v))}{\Lambda(\mathcal{R}_\eta(v))} \right) = \int_{J_0^\eta} E_{\theta, \lambda} \left(\frac{\mathbf{1}_{\{p_\eta(x) \leq v\}}}{\Lambda(\mathcal{R}_\eta(v))} \right) d\Lambda(x).$$

Then, choosing v such that

$$\frac{\Lambda(\mathcal{R}_\eta(v))}{\Lambda(\mathcal{X}_\eta)} \geq \frac{v}{\alpha}$$

and doing as if v was deterministic leads to

$$\text{FDR}_{\theta, \lambda}^\eta(v) \leq \alpha \int_{J_0^\eta} v^{-1} E_{\theta, \lambda} \left(\frac{\mathbf{1}_{\{p_\eta(x) \leq v\}}}{\Lambda(\mathcal{X}_\eta)} \right) d\Lambda(x),$$

which leads to

$$\text{FDR}_{\theta, \lambda}^\eta(v) \leq \frac{\alpha}{\Lambda(\mathcal{X}_\eta)} \int_{J_0^\eta} v^{-1} P_{\theta, \lambda}(p_\eta(x) \leq v) d\Lambda(x) \leq \alpha \Lambda(J_0^\eta) / \Lambda(\mathcal{X}_\eta) \leq \alpha,$$

by using the property of p -values (4). The above heuristic suggests to choose V^α by using

$$V^\alpha = \max \left\{ v \geq 0 : \frac{\Lambda(\mathcal{R}_\eta(v))}{\Lambda(\mathcal{X}_\eta)} \geq v/\alpha \right\}.$$

Weighted-BH procedure. Quantile V^α is random (even conditionally on \mathfrak{N}) because it depends on the observed p -value process. Nevertheless, since the p -value process is piece-wise constant in our framework, finding threshold V^α is very simple. We propose a step-up algorithm that is inspired from the famous *weighted BH threshold* (Benjamini and Hochberg, 1997), with the essential difference that our algorithm is based on random weights. Recall that the processes $(S_\eta(x))_x$ and therefore $(p_\eta(x))_x$ are piece-wise constant on the random partition $\tau = (\tau_m)_{0 \leq m \leq M}$ ($\tau_0 = 0$, $\tau_M = 1$). Therefore

$$\Lambda(\mathcal{R}_\eta(v)) = \sum_{m=1}^M (\tau_m - \tau_{m-1}) \mathbf{1}_{\{p_\eta(\tau_{m-1}) \leq v\}}.$$

This implies that the threshold V^α defined in (3.4) can be derived as follows:

- Compute the weights,

$$w_m = \frac{\tau_m - \tau_{m-1}}{\Lambda(\mathcal{X}_\eta)}, \quad , 1 \leq m \leq M,$$

- Denote $p_m = p_\eta(\tau_{m-1})$, $1 \leq m \leq M$,
and order these p -values in increasing order $p_{\sigma(1)} \leq \dots \leq p_{\sigma(M)}$ for an appropriate permutation σ of $\{1, \dots, M\}$;
- Consider $\hat{k} = \max\{k \in \{1, \dots, M\} : p_{\sigma(k)} \leq \alpha \sum_{\ell=1}^k w_{\sigma(\ell)}\}$ (with $\hat{k} = 0$ if the set is empty);
- Compute V^α as $\alpha \sum_{\ell=1}^{\hat{k}} w_{\sigma(\ell)}$.
- Define $\mathcal{R}_\eta^{wBH} = \{x \in \mathcal{X}_\eta : p_\eta(x) \leq V^\alpha\}$.

By contrast with the classical BH step-up procedure, each p -value has its own weight w_m which is not uniform but adaptive to the data. This weight means that the quantity of interest is not the number of tests, i.e. the cardinality of the partition τ , but the proportion of "time" that the p -value process spends at a particular value.

Remark 3.4. *It is worth to note that $\mathcal{R}_\eta^{wBH} = \{x \in \mathcal{X}_\eta : q_\eta(x) \leq \alpha\}$ where $q_\eta(x)$ stands for the adjusted p -value*

$$q_\eta(x) = \min_{k: p_{\sigma(k)} \geq p_\eta(x)} \left\{ \frac{p_{\sigma(k)}}{\sum_{\ell=1}^k w_{\sigma(\ell)}} \right\}.$$

Theoretical result. Our procedure is based on a heuristic that has been made fully theoretically valid by Blanchard et al. (2014) under suitable assumptions of measurability and positive dependence (namely, finite dimensional positively regressively dependent p -value process). The difficulty comes from the interplay between the numerator and the denominator within the expectation in (3.4) via v ,

because $v = V^\alpha$ is a random variable. While the measurability issue is not critical here (by the piece-wise constant property), the positive dependence condition seems difficult to check in general. For instance, this condition is known to exclude the two-sided test statistic (Yekutieli, 2008). We prove here that this positive dependence condition holds in the one-sided testing setting (both for the homogeneity test with a known λ and the two-sample test), by using the peculiar properties of Poisson processes.

Theorem 3.5. *For all $\alpha \in (0, 1)$, for a p -value process given either by Eq. 1 or Eq. 2, for one-sided null hypotheses, the corresponding procedure \mathcal{R}_η^{wBH} , satisfies*

$$\forall(\theta, \lambda), \text{FDR}_{\theta, \lambda}^\eta(\mathcal{R}_\eta^{wBH}) \leq \alpha.$$

This result is proved in the Supplementary File. Theorem 3.5 hence theoretically justifies the use of a weighted step-up procedure to control the continuous FDR. However, it is only valid for simple versions of our test statistic, which ensure the positive dependence structure. For more general cases, Blanchard et al. (2014) proposed a modification of the step-up algorithm. This modification consists in introducing a function $\beta(v) \leq v$ of a specific form, such that the quantity $\sum_{\ell=1}^k w_{\sigma(\ell)}$ is replaced by the smaller quantity $\beta\left(\sum_{\ell=1}^k w_{\sigma(\ell)}\right)$. However, this modification is reported to be quite conservative (Blanchard and Roquain, 2008). Furthermore, our simulation results show that the FDR control (3.4) is achieved in most of the situations we studied, even for more complex p -value processes than the ones we considered here (homogeneity test with known intensity, and two-sample test)

Note that the homogeneity test with known intensity in Theorem 3.5 is similar to Corollary 4.5 in Blanchard et al. (2014). However, a subtle distinction concerns the way the null hypotheses are defined (windows versus points). A consequence is that our approach does not rely on any knowledge regarding the regularity of the underlying process intensity.

4 Kernel-based statistics for sliding windows

4.1 On the limitations of the count-based statistic

Let us focus on the two-sample test to illustrate our motivations. In this context, our objective is to detect local differences between unknown functions ν_A and ν_B based on individual tests defined on a set of windows. In a first approximation, this can be achieved by comparing $N_A(I)$ with $N_B(I)$, the count of occurrences on window I , that are distributed as a couple of independent Poisson variables. However, given that $N_A(I)$ (resp. $N_B(I)$) is an approximation of $\int_I \nu_A(t)dt$ (resp. $\int_I \nu_B(t)dt$), the count based-statistic strategy can be considered as an hypothesis that tests the equality of function integrals rather than the equality of the functions themselves. Hence the induced testing procedure may lack of accuracy.

Therefore we developed a kernel-based strategy to propose single test statistics that do not depend on the sole number of occurrences in a window, but that consider its whole composition as defined in Section 2.5. The interest of kernel-based statistics is that they account the physical distances between occurrences, which explains the gain of performance.

4.2 The kernel vision in the two-sample case

The kernel framework for testing proposed by Gretton et al. (2012) and Fromont et al. (2013) has been especially designed to separate functions, and not just their integrals. It is based on the estimation of a distance between the functions ν_A and ν_B (in a RKHS, or in L^2). In particular, Fromont et al. (2013) show that these test statistics are minimax with respect to various smoothness classes, and in this sense, powerful against some alternatives that cannot be distinguished by count-based tests.

In the following we consider translation invariant kernels denoted by K , and in practice we will restrict ourselves (see Sections 6 and 7) to the Gaussian kernel that has been shown to be the most powerful on simulations (Fromont et al., 2013). The shape of the statistic is then given by

$$S_\eta(x) = \sum_{T \neq T' \in N \cap I_\eta(x)} K(T - T') \varepsilon_T \varepsilon_{T'},$$

where N is the joint process and the ε_T 's are the marks (see Section 2.1). In particular, it only depends on the composition of the window. To understand, why this statistic is meaningful, let us compute its expectation. By classical computations for Poisson processes,

$$\begin{aligned} E_{\theta, \lambda}(S_\eta(x)) &= E_\lambda \left[\sum_{T \neq T' \in N \cap I_\eta(x)} K(T - T') \theta(T) \theta(T') \right] \\ &= \int_{I_\eta(x)^2} K(t - t') \theta(t) \theta(t') \lambda(t) \lambda(t') dt dt' \\ &= \int_{I_\eta(x)^2} K(t - t') [\nu_A(t) - \nu_B(t)] [\nu_A(t') - \nu_B(t')] dt dt' \end{aligned}$$

If $K \star f$ denote the convolution product, we end up with

$$E_{\theta, \lambda}(S_\eta(x)) = \int_{I_\eta(x)} (K \star [(\nu_A - \nu_B) \mathbf{1}_{I_\eta(x)}]) (t) [\nu_A(t) - \nu_B(t)] dt.$$

In particular if K is a convolution kernel with bandwidth h we obtain, under classical assumptions, that

$$E_{\theta, \lambda}(S_\eta(x)) \xrightarrow{h \rightarrow 0} \int_{I_\eta(x)} [\nu_A(t) - \nu_B(t)]^2 dt.$$

Procedure	test statistic	computations	homogeneity test	two-sample test
min- p	count	p_η full procedure	exact	1-step MC
	kernel	p_η full procedure	1-step MC	2-step MC
wBH	count	p_η full procedure	exact	exact
	kernel	p_η full procedure	1-step MC	1-step MC

Table 1: Summary of the required computations to implement the continuous tests in practice for the homogeneity and two-sample test.

This short derivation illustrates the motivation to use such a kernel-based statistic, as a good estimate of a distance between intensity functions. The one-sided version of this statistic, as well as the homogeneity case are presented in the Supplementary Material. Note that from a computational point of view, the complexity for computing $S_\eta(x)$ is in $O(N^2(I_\eta(x)))$. However, thanks to the piece-wise constant property of $S_\eta(x)$ on partition τ , the continuum $\{S_\eta(x), x \in \mathcal{X}_\eta\}$, can be computed using the in/out points of each scanning window, which drastically reduces the global computational burden.

5 Monte-carlo procedures for p -values estimation and adjustment

5.1 The distribution of the single test statistic being explicit.

For tests based on the count statistic the p -value process is explicitly computable. Hence, no need for Monte-Carlo in this case. Moreover, since our weighted BH procedure (Section 3.2) is completely explicit, \mathcal{R}_η^{wBH} can be easily computed in practice. In this situation we have already proved the FDR control for some particular cases and we will show that it holds in more general cases by using simulations (Section 6).

However, the min- p procedure is not explicit since it requires the conditional distribution of $\inf_{x \in \mathcal{X}_\eta} p_\eta(x)$ given \mathfrak{N} under the full null hypothesis (Section 3.2). Therefore we need to use a Monte-Carlo procedure to approach this distribution. This can be achieved while maintaining a valid FWER control by following Romano and Wolf (2005a), as detailed in Section 5.4 below.

5.2 The distribution of the single test statistic being not explicit

This concerns kernel-based tests (Section 4) for which the p -value process is not even computable. One can use again the estimated p -values developed by Romano

and Wolf (2005b) since their result shows that they satisfy (4) and \mathcal{R}_η^{wBH} can be easily computed in practice. However, there is no theoretical FDR guarantee and we will only show the FDR control by simulations in Section 6. Consequently, to compute $\hat{p}_\eta(x)$, the weighted-BH procedure only requires one Monte-Carlo step in this situation.

As for the min- p procedure (Section 3.2), it requires an additional Monte-Carlo step. To this end, a classical and sufficient method is to perform an approximation using two independent Monte-Carlo samples. We detail this approach for the Homogeneity case in the supplementary material. For the two-sample case, the Monte-Carlo procedure is quite demanding and is close in spirit to bootstrap methods. In this case, we developed a double Monte-Carlo procedure, which only requires *one* Monte-Carlo sample and still guarantees a controlled FWER (Section 5.4).

5.3 Estimation of the p -value process by Monte-Carlo

Recall that this estimation is required when using test statistics different from the statistics based on counts. The full procedures are summarized in Table 1.

Homogeneity test. Here, the conditioning variable $\mathfrak{N} = N([0, 1])$ is the same for all scanning windows and the distribution of $S_\eta(x)$ under $\mathcal{H}_0\{I_\eta(x)\}$ does not depend on x . Let us focus on the first window $I_\eta(\eta/2) = (0, \eta]$ and simulate points inside this window under $\mathcal{H}_0\{I_\eta(\eta/2)\}$. Hence we consider B random samples of the number of points $N^b(I_\eta(\eta/2))$ such that

$$\forall b \in \{1, \dots, B\}, N^b(I_\eta(\eta/2)) | \{N([0, 1]) = n\} \underset{iid}{\sim} \mathcal{B}(n, \eta),$$

and then we draw

$$(T_1, \dots, T_{N^b(I_\eta(\eta/2))}) \underset{iid}{\sim} U((0, \eta]),$$

to sample positions of a homogeneous Poisson process N^b inside window $(0, \eta]$. For each b , one can compute thanks to the positions of the points inside $[0, \eta]$ the value of the test statistic for the first window, that we denote S_η^b . But since all the statistics are distribution invariant by translation under the null, one can also say that this sample has the same distribution as $S_\eta(x)$ under $\mathcal{H}_0\{I_\eta(x)\}$ for every window center $x \in \mathcal{X}_\eta$.

Therefore, whatever B , one has easily access to $(S_\eta^1, \dots, S_\eta^B)$, B i.i.d. variables with the same distribution as $S_\eta(x)$ under $\mathcal{H}_0\{I_\eta(x)\}$ for every x . Then, following Romano and Wolf (2005b), we define:

$$\hat{p}_\eta(x) = \frac{1}{B+1} \left(1 + \sum_{b=1}^B \mathbf{1}_{\{S_\eta^b \geq S_\eta^0(x)\}} \right),$$

which amounts to use the empirical function $F_{\mathfrak{N},x}$ defined in Eq. 3 over a $B+1$ sample where $S_\eta^0(x)$ is the observed statistic.

Two sample test. We label the observed marks $\varepsilon^0 := \varepsilon = (\varepsilon_T)_{T \in N}$, such that it constitutes the first term of a $(B + 1)$ -sample of marks, filled by B independent draws of *i.i.d.* Rademacher sets, $\varepsilon^b := (\varepsilon_T^b)_{T \in N}$, for $b = 1, \dots, B$. The conditional distribution given $N = \mathfrak{N}$ of these Rademacher processes are:

$$\varepsilon_T^b | N \underset{iid}{\sim} 2\mathcal{B}\left(\frac{1}{2}\right) - 1,$$

that is the distribution that the observed marks ε_T 's should have under the full null hypothesis. As previously, the Monte-Carlo approximated conditional p -value of window $I_\eta(x)$ is defined by:

$$\widehat{p}_\eta(x) = \frac{1}{B + 1} \left(1 + \sum_{b=1}^B \mathbf{1}_{\{S_\eta^b(x) \geq S_\eta^0(x)\}} \right), \quad (10)$$

where $S_\eta^0(x)$ is the observed statistic and where $S_\eta^b(x)$ is the value of the test statistic on $I_\eta(x)$ that is computed with the resampled marks ε^b and fixed joint process N . Let us underline that because in the two-sample case the partition τ only depend on N and do not vary with the resampling scheme ε^b , the processes $(S_\eta^b(x))_{x \in \mathcal{X}_\eta}$, for $b = 0, \dots, B$, and therefore the process $(\widehat{p}_\eta(x))_{x \in \mathcal{X}_\eta}$ are piece-wise constant on τ .

5.4 Two-step Monte-Carlo method for the min- p procedure.

Recall that when using the min- p procedure with a p -value process that is not explicit, we need two Monte-Carlo approximations.

Homogeneity test. In that case, the distribution of the test statistics under the null are independent of the window centers, while the distribution of the minimum of the p -values relies on the observations over all windows. Hence, it is quite natural to perform two separated Monte-Carlo schemes for these two steps.

Two sample test. In the two-sample case, the trick is that one can use the *same sample* of (N, ε^b) 's as the one used in (10) for both approximations. More precisely, the p -values $\widehat{p}_\eta^b(x)$, corresponding to the simulated process (N, ε^b) , can be computed as

$$\forall b \in \{0, \dots, B\}, \forall x \in \mathcal{X}_\eta, \widehat{p}_\eta^b(x) = \frac{1}{B + 1} \sum_{b'=0}^B \mathbf{1}_{\{S_\eta^{b'}(x) \geq S_\eta^b(x)\}},$$

that is we do the same computation as in Eq. 10 except that we replace the observed statistics $S_\eta^0(x)$ by the simulated one $S_\eta^b(x)$, hereby computing a "resampled" p -value process corresponding to ε^b . Then, let us define, for all $b = 0, \dots, B$

$$m^b = \inf_{x \in \mathcal{X}_\eta} \widehat{p}_\eta^b(x).$$

The corresponding adjusted Monte-Carlo p -values are then given by

$$\hat{q}_\eta(x) = \frac{1}{B+1} \left(1 + \sum_{b=1}^B \mathbf{1}_{\{m^b \leq \hat{p}_\eta(x)\}} \right).$$

Finally the rejection set is given by

$$\hat{\mathcal{R}}_\eta^m = \{x \in \mathcal{X}_\eta, \hat{q}_\eta(x) \leq \alpha\}.$$

The previous double Monte-Carlo procedure is based on the use of the same random signs $(\varepsilon_T^b)_{T \in N}$, $0 \leq b \leq B$ and not on two sets of different simulations. Indeed the ε^b 's are used twice: a first time to estimate the p -values of the observed process and a second time to give the p -values even on the simulated realizations leading to the adjusted p -values. This procedure has the double advantage of sparing computational resource, while enjoying the same control property as shown hereafter.

Theorem 5.1. *Let $\alpha \in (0, 1)$. For the two-sample case, the continuous testing procedure defined by (5.4) satisfies*

$$\forall(\theta, \lambda), FWER_{\theta, \lambda}^\eta(\hat{\mathcal{R}}_\eta^m) \leq \alpha.$$

This result holds whatever the choice of the single test statistic that only depends on the composition of the window.

6 Simulations

6.1 Homogeneity test

For a fixed parameter r , we simulate a homogeneous Poisson process corresponding to the following piecewise constant signal θ (illustrated in Figure 9, Supp. Material):

$$\theta(t) = \begin{cases} +\theta^*, & \forall t \in \mathcal{I}_1^+, \\ -\theta^*, & \forall t \in \mathcal{I}_1^-, \\ 0, & \text{otherwise.} \end{cases}$$

with

$$\begin{cases} \mathcal{I}_1^+ = [\frac{1}{4} - \frac{r}{4}, \frac{1}{4} + \frac{r}{4}] \cup [\frac{1}{2} - \frac{r}{4}, \frac{1}{2} + \frac{r}{4}] \cup [\frac{3}{4} - \frac{r}{4}, \frac{3}{4} + \frac{r}{4}], \\ \mathcal{I}_1^- = [\frac{1}{4} - \frac{r}{2}, \frac{1}{4} - \frac{r}{4}] \cup [\frac{1}{4} + \frac{r}{4}, \frac{1}{4} + \frac{r}{2}] \\ \quad \cup [\frac{1}{2} - \frac{r}{2}, \frac{1}{2} - \frac{r}{4}] \cup [\frac{1}{2} + \frac{r}{4}, \frac{1}{2} + \frac{r}{2}] \\ \quad \cup [\frac{3}{4} - \frac{r}{2}, \frac{3}{4} - \frac{r}{4}] \cup [\frac{3}{4} + \frac{r}{4}, \frac{3}{4} + \frac{r}{2}]. \end{cases}$$

and $\mathcal{I}_1 = \mathcal{I}_1^+ \cup \mathcal{I}_1^-$, $\mathcal{I}_0 = [0, 1] \setminus \mathcal{I}_1$, the positions on which $\theta(\cdot)$ is not null (or null, respectively). We consider this particular shape of θ to ensure that $\int_0^1 \theta(t) dt = 0$.

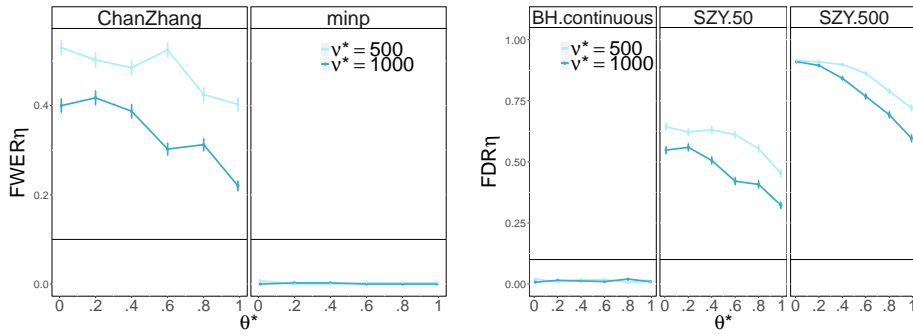


Figure 2: Comparison of asymptotic vs. non asymptotic methods for the control of Type-I error rates ($\alpha = 0.1$). ChanZhang: asymptotic approximation of the scan statistics distribution proposed by Chan and Zhang (2007), SYZ: approximation of the FDR proposed by Siegmund et al. (2011) with 50 (SZY.50) and 500 (SZY.500) windows of identical length. FWER_η and FDR_η refer to Eqs. 5 and 6.

For a fixed background intensity parameter ν^* , we therefore simulate 1000 times a Poisson process N of intensity $\nu(\cdot) = \nu^*(1 + \theta(\cdot))$. Parameter θ^* is used to tune the distance to homogeneity (which reflects the difficulty of the test), with $\theta^* = 0.01$ standing for a situation where almost every position is under the homogeneous regime, $\theta^* = 0.99$ standing for a situation with high separability. Parameter ν^* controls the number of occurrence and varies in $\{500, 1000\}$. To compare with other scanning procedures, we opt for one-sided tests (the alternative corresponding to too large θ), with fixed window size $\eta = 2r$. More precisely, we compare our count-based min- p and count-based weighted BH procedures in continuous time with asymptotic approaches based on the scan statistic that controls the FWER (Chan and Zhang, 2007) and the FDR (Siegmund et al., 2011).

We consider this homogeneity test to compare asymptotic approaches based on the scan statistic to control the FWER (Chan and Zhang, 2007) and the FDR (Siegmund et al., 2011), with our min- p and weighted BH procedures in continuous time, based on the count statistic. Figure 2 clearly shows that our min- p procedure in continuous time controls the FWER even for moderate sample sizes, whereas the scan statistic procedure requires a strong signal ($\theta^* \sim 1$), or a higher number of occurrences ($\nu^* = 5000$, not shown). This trend is even more clear for the FDR control, that is not ensured by method of Siegmund et al. (2011). Note that contrary to continuous testing, this method requires a discretization step on which the performance depend. In this case we choose fixed-length non-overlapping windows to compute their procedure.

6.2 Two sample test

To test the differences between intensities ν_A and ν_B , we sample a joint point process N over $[0, 1]$ with constant intensity $\nu(\cdot) = \nu^* \in \{500, 1000\}$. Then




Figure 3: Comparison of the count and the Gaussian kernel statistics for the control of joint error rates ($\alpha = 0.1$).

we sample the labels ε^0 given the occurrences of the joint process $T \in N$, in an independent manner, such that

$$\forall T \in N, \varepsilon_T^0 | N \sim \begin{cases} 2\mathcal{B}((\theta^* + 1)/2) - 1 & \text{if } T \in \mathcal{I}_1 \\ 2\mathcal{B}(1/2) - 1 & \text{if } T \in \mathcal{I}_0 \end{cases}$$

Thus positions $T \in N$ with $\varepsilon_T^0 = +1$ (resp. -1) stand for points of process N_A (resp. N_B) with intensity ν_A (resp. ν_B). Parameter θ^* represents the distance between ν_A and ν_B (which reflects the difficulty of the test), with $(\theta^* + 1)/2 = 0.51$ standing for a situation where the two intensities are nearly equal, and $(\theta^* + 1)/2 = 0.9$ standing for a situation with high separability. \mathcal{I}_0 and \mathcal{I}_1 are defined as in the homogeneity case, and each configuration is repeated 1000 times. For each value of θ^* , bootstrap p -values are computed using $B = 10^5$ Monte Carlo samples, along with the Gaussian kernel K of bandwidth $h = \eta$. Decision rules are based either on the min- p procedure, or on the weighted BH procedure, and we are comparing the performance of the count statistic (one-sided) with the statistic based on the Gaussian Kernel (one-sided). The size of the windows is fixed $\eta = 2r$.

Even if the FWER and the FDR are well controlled by our two procedures, the Gaussian kernel statistic seems to be less conservative when controlling the FWER and seems more robust changes in the intensity ν^* of the process (Fig. 3). The control of the FDR are comparable between the two statistics. However, when the kernel and the count statistics are compared using sensitivity and specificity (Fig. 4), for the same level of signal (same θ^*), the kernel statistics is globally more accurate than the count statistics (higher sensitivity and specificity). This can be explained by the fact that kernel-based statistics are in fact U -statistics, hereby using information from pairs of points through the kernel. The difference of these statistics will also be illustrated on the analysis of experimental data.




Figure 4: ROC Curve to compare the Count statistic (dashed lines) with the Gaussian Kernel statistics (plain lines). Sensitivity, or true positive rate (proportion of true positive windows among all rejected windows). Specificity, or true negative rate, proportion of true negative windows among all accepted windows. One dot corresponds to a value for α . Here $\nu^* = 1000$, and p -values were adjusted using the wBH procedure.

7 Applications

Description of the full procedure in practice. In the following applications, we compute the p -value process as well as its min- p and wBH adjusted versions. In order to propose a completely operational procedure, we need to clarify the distinction between the points space and the windows space on which the decision is made. Indeed, if window $I_\eta(x)$ is accepted by our procedure, and if there is no mistake, then one is sure that $H_{0,t}$ holds for all point t in $I_\eta(x)$, whereas if a window $I_\eta(x)$ is rejected without mistake, it is yet not clear whether there is not a smaller interval included in $I_\eta(x)$ on which the null hypothesis holds. Consequently we start by defining the set of accepted windows

$$\widehat{\mathcal{A}}_\eta = \{x \in \mathcal{X}_\eta : \widehat{q}_\eta(x) > \alpha\},$$

with $\widehat{q}_\eta(x)$ denoting the estimated adjusted p -value process. Finally, the graphical representation of our procedure is based on the estimated set of points that lie within rejected windows, namely: $\widehat{\mathcal{I}}_1^\eta = [0, 1] \setminus \widehat{\mathcal{I}}_0^\eta$, where

$$\widehat{\mathcal{I}}_0^\eta = \left\{ t \in [0, 1] : \exists x \in \widehat{\mathcal{A}}_\eta \text{ s.t. } t \in I_\eta(x) \right\}.$$

7.1 Do neurons respond to some stimulus ?

In Neurosciences, the succession of action potentials of a given neuron, also called spike trains, constitutes the key biological signal that carries information regarding

the functioning of neurons. In particular, understanding the rapid changes in firing rates due to a certain stimulus (Kass et al., 2003) is a challenging task, and inhomogeneous Poisson processes are classically used to model such data (Pipa et al., 2013). Asking whether there is a change of firing rate after exposure typically refers to our homogeneity test, and we are seeking for time intervals on which the rate is modified with respect to the standard behaviour. To illustrate our procedure we analyzed public experimental data (Pouzat and Chaffiol, 2009) that consist in spike trains of a cockroach neuron stimulated by different odors like citronellal or terpineol. A first question is to determine if this neuron responds to the stimulus (one-sided homogeneity test), and if the response is different from one stimulus to another (two-sided two-sample test).

The neuronal activity starts to be recorded at time 0, and the citronellal is presented between 0.42 and 0.45 (normalized time, Fig. 5). Using the count statistic and the $\min-p$ procedure, our method identifies a neuronal response only when the neuron is stimulated (Fig. 5, dark-grey rectangle), but not after the stimulus, whereas the weighted BH procedure identifies two firing rate changes posterior to the stimulus. Interestingly, the weighted BH procedure also identifies a deviation before the stimulus, which may correspond to habituation and therefore anticipation of the stimulus by the animal (see for a similar phenomenon with synchrony instead of firing rates modifications (Tuleau-Malot et al., 2014)). What is remarkable is that the Gauss Kernel statistic globally identifies the same intervals of changes when compared with the count statistic (Fig. 6), with the notable difference that the instantaneous response to the stimulus is located more precisely, which illustrates the higher specificity of the kernel-based statistics. Moreover, the kernel detects an additional interval of change in the firing rate. When studying the differential response of the same neuron to different stimuli (citronellal, terpineol) the differential change of fire rate occurs just after the stimulus (Fig. 7). This difference is identified by the weighted BH procedure, but not by the $\min-p$, which lacks of power, as already documented.

7.2 Differential analysis of replication origins between cell lines

Next Generation Sequencing technologies (NGS) are based on the massive parallel sequencing of short DNA fragments, which have given rise to a new area in the field of high throughput biology. The applications of such technologies are many, and we focus on the particular cases where NGS allow the fine mapping of genomic features along the genome, like transcription factors (using the ChIP-Seq technology), or copy number variations (CNV-Seq) for instance. When applied to the field of replication, NGS technologies have allowed the fine-scale mapping of replication origins along the human genome. Briefly, replication is the mechanism by which genomes are duplicated into two exact copies. Genomic stability is under the control of a spatio-temporal program that orchestrates both the positioning and the timing of firing 100,000 replication starting points, also called replication origins (Picard et al., 2014). This fine mapping has allowed us to identify potential




Figure 5: Test of homogeneity for the spike trains of a cockroach neuron (neuron 2) excited by citronellal at $[0.42, 0.45]$ (normalized time, red band). Top: count statistic $S_\eta(t)$ with $\eta = 1/20$. Light-grey rectangles stand for intervals $\hat{\mathcal{I}}_1^\eta$ on which the homogeneity hypothesis is rejected by the wBH procedure. The dark grey rectangle corresponds to the interval $\hat{\mathcal{I}}_1^\eta$ that is rejected by the min- p procedure ($\alpha = 5\%$). Bottom: p -value process adjusted with the min- p or the wBH procedures.




Figure 6: Test of homogeneity ($\eta = 1/20$) for the spike trains of a cockroach neuron (neuron 2) excited by citronellal at $[0.42, 0.45]$ (normalized time, red vertical arrows). Rectangles stand for intervals \hat{I}_1^η on which the homogeneity hypothesis is rejected by the wBH procedure ($\alpha = 5\%$).

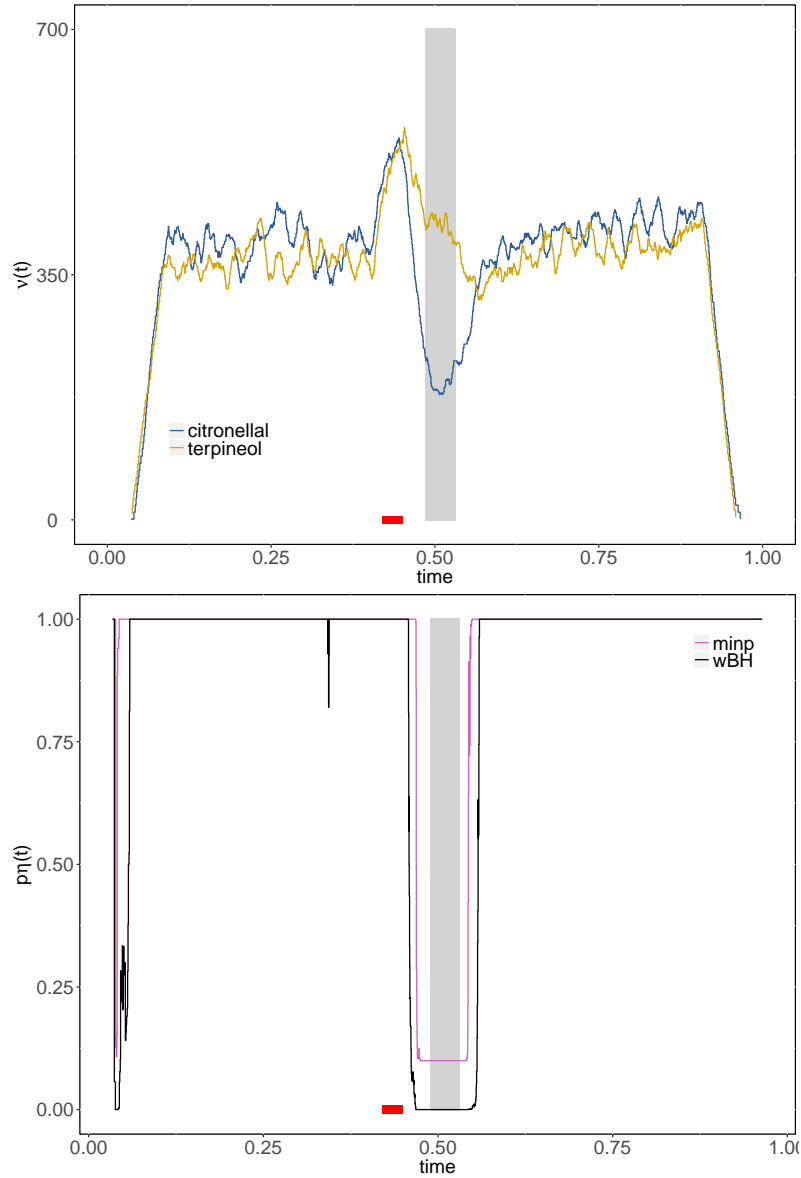


Figure 7: Two-sample test for the spike trains of a cockroach neuron (neuron 2) excited by citronellal at $[0.42, 0.45]$ (normalized time, red vertical arrows). Top: count statistic $S_\eta(t)$ with $\eta = 1/20$. The dark grey rectangle corresponds to the interval $\widehat{\mathcal{I}}_1^\eta$ that is rejected by the wBH procedure ($\alpha = 5\%$). Bottom: p -value process adjusted with the min- p or the wBH procedures.

regulators of the replication dynamics, and the comparative analysis of replication origins maps can help to identify a core of origins that are common to all human cells, compared with cell-type specific origins more related to cell differentiation. Moreover, pathologies like cancer being linked to replication instabilities, comparing replication origins in cancer cells vs. non-cancer cells could be promising.

The chromosomal organization of replication origins can easily be modeled by a heterogeneous Poisson process of unknown intensity, and we consider the comparison of origins between cancer cells (Hela) and embryonic stem cells (iPS). Since the total number of origins differ from one condition to another and that globally there is no real difference in shape between both intensities except at some particular positions of interests, we need to account for the difference in the total number of occurrences per condition. To proceed, we modify the conditional distribution of the marks under the null, which becomes conditional to the total (random) number of points coming from condition A and B respectively:

$$\varepsilon_T | N, \lambda_A, \lambda_B \underset{\mathcal{H}_0}{\sim} 2\mathcal{B}\left(\frac{\lambda_A}{\lambda_A + \lambda_B}\right) - 1,$$

with $\lambda_A = \int_0^1 \nu_A(t) dt$ and $\lambda_B = \int_0^1 \nu_B(t) dt$. Since in practice λ_A and λ_B are unknown, they are replaced by the total number of points in N_A and N_B . This procedure resumes to testing the equality of density functions instead of intensities.

As expected, origins densities show strong variations along chromosomes, but those variations are globally comparable between cell lines (Fig. 8). One region shows significantly more origins in Hela cells (cancer cells), that is detected by both the count and the Gaussian kernel statistic after the continuous wBH procedure. The min- p adjustment does not allow the detection of this region at the same nominal level, but the variations of the adjusted p -value process indicate a strong difference of densities between cell lines. This region corresponds to a gene-rich region that needs investigation for further biological characterization. More and more cell lines have been investigated to determine the core of replication origins that are common to any human cells, and the set of origins that are cell-type specific (Besnard et al., 2012). This context offers a nice perspective of the present work, by generalized our two-sample test to the more general framework of a k -sample (or ANOVA-like) test.

8 Conclusion

We proposed a new framework for testing Poisson processes intensities based on continuous testing. This framework appears very powerful, from both the theoretical and the practical points of views. It allows the complete definition of the testing procedure, based on local null hypothesis, which provides a clear definition of false positive windows, that was lacking in existing works. We introduced the p -value process that becomes a stochastic process that needs to be adjusted to account for joint error rates. Contrary to existing approaches, continuous testing avoid the




Figure 8: Top: density of replication origins along chromosome 16-q, for HeLa and iPS cells. The grey rectangle corresponds to the interval $([71.64, 71.91])$ in Megabases) on which the equality hypothesis is rejected at level $\alpha = 0.1$ using the count statistic, after adjustment by the wBH procedure. The Gaussian kernel statistics gives the same result. Bottom: p -value process adjusted by the wBH procedure (black), and the min- p procedure (purple). The min- p procedure does not reject any region. The size of the window corresponds to windows of 1 Mb.

arbitrary discretization of this process, which is coherent with the continuous nature of the underlying model. The continuous framework especially provides a rigorous definition of error rates in continuous time, and we proposed two procedures to control the FWER and the FDR. In addition to theoretical controls, we show the empirical properties of our procedure, that can be easily implemented, and is available in the form of the `contest` R-package. The application on experimental data inspires extensions of our framework, that could be adapted to the testing of more than 2 conditions, or to other tests like local independence tests that are essential in Neurosciences (Albert et al., 2015). Calibrating the window size is certainly one of the major perspective of this work. While the subject has been extensively investigated in the context of model estimation, only a few works considered the question in the testing framework (Gao and Gijbels, 2008). Consequently, one challenge ahead will be to determine if the continuous testing framework would be appropriate to derive a calibration method to automatically select the resolution of the test from the data. Other future extensions involve either spatial aspects or more general point processes from a theoretical point of view.

Acknowledgment

We are grateful to Anne-Laure Fougères for fruitful discussions on the subject, to Matthieu Lerasle for sharing with us the trick of the double Monte-Carlo step, and to Ivan Bardet for his work on the Chan and Zhang (2007) paper as an intern. This research was partly supported by the french Agence Nationale de la Recherche (ANR 2011-BS01-010-01 Calibration project, ANR-11-BINF-0001-06 ABS4NGS project, ANR-15-CE12-0004-02 OriMolMech project and ANR-16-CE40-0019 SansSouci project). This work was performed using the computing facilities of the computing center LBBE/PRABI.

References

- Albert, M., Y. Bouret, M. Fromont, and P. Reynaud-Bouret (2015, 12). Bootstrap and permutation tests of independence for point processes. *Ann. Statist.* 43(6), 2537–2564.
- Arias-Castro, E., R. M. Castro, E. Tánchez, and M. Wang (2015, August). Distribution-Free Detection of Structured Anomalies: Permutation and Rank-Based Scans. *ArXiv e-prints*.
- Benjamini, Y. and Y. Hochberg (1995). Controlling the false discovery rate: a practical and powerful approach to multiple testing. *J. Roy. Statist. Soc. Ser. B* 57(1), 289–300.
- Benjamini, Y. and Y. Hochberg (1997). Multiple hypotheses testing with weights. *Scand. J. Statist.* 24(3), 407–418.
- Besnard, E., A. Babled, L. Lapasset, O. Milhavet, H. Parrinello, C. Dantec, J. M. Marin, and J. M. Lemaitre (2012, Aug). Unraveling cell type-specific and reprogrammable human replication origin signatures associated with G-quadruplex consensus motifs. *Nat. Struct. Mol. Biol.* 19(8), 837–844.
- Blanchard, G., S. Delattre, and E. Roquain (2014). Testing over a continuum of null hypotheses with False Discovery Rate control. *Bernoulli* 20(1), 304–333.
- Blanchard, G. and E. Roquain (2008). Two simple sufficient conditions for FDR control. *Electron. J. Stat.* 2, 963–992.

- Chan, H. and N. Zhang (2007). Scan statistics with weighted observations. *Journal of the American Statistical Association* 102(478), 595–602.
- Daley, D. J. and D. Vere-Jones (2008). *An introduction to the theory of point processes. Vol. I-II* (Second ed.). Probability and its Applications (New York). Springer.
- Dudoit, S., J. P. Shaffer, and J. C. Boldrick (2003, 02). Multiple hypothesis testing in microarray experiments. *Statist. Sci.* 18(1), 71–103.
- Dümbgen, L. and V. G. Spokoiny (2001, 02). Multiscale testing of qualitative hypotheses. *Ann. Statist.* 29(1), 124–152.
- Fromont, M., B. Laurent, and P. Reynaud-Bouret (2013). The two-sample problem for poisson processes: adaptive tests with a non-asymptotic wild bootstrap approach. *Annals of Statistics* 41(3), 1431–1461.
- Fromont, M., M. Lerasle, and P. Reynaud-Bouret (2016). Family wise separation rate for multiple testing. *to appear in the Annals of Statistics*.
- Fu, J. C., T.-L. Wu, and W. Y. W. Lou (2012). Continuous, discrete, and conditional scan statistics. *J. Appl. Probab.* 49(1), 199–209.
- Gao, J. and I. Gijbels (2008). Bandwidth selection in nonparametric kernel testing. *Journal of the American Statistical Association* 103(484), 1584–1594.
- Goeman, J. J. and A. Solari (2010). The sequential rejection principle of familywise error control. *The Annals of Statistics* 38(6), 3782–3810.
- Gretton, A., K. Borgwardt, M. Rasch, B. Schoelkopf, and A. Smola (2012). A kernel two-sample test. *Journal of Machine Learning Research* 13, 723–773.
- Grün, S. (1996). *Unitary joint-events in multiple-neuron spiking activity: Detection, significance and interpretation*. Ph. D. thesis, Thun: Verlag Harri Deutsch.
- Kass, R., V. Ventura, and C. Cai (2003). Statistical smoothing of neuronal data. *Network: Comput. Neural Syst.* 14, 5–15.
- Kingman, J. (1993). *Poisson Processes*. Oxford Studies in Probability.
- Kulldorf, M. (1997). A spatial scan statistics. *Communications in statistics: Theory and methods* 26, 1481–1492.
- Kulldorf, M., R. Heffernan, J. Hartman, R. Assunção, and F. Mostashari (2005). A space–time permutation scan statistic for disease outbreak detection. *PLoS Medicine* 2(3), e59.
- Kulldorf, M., L. Huang, and K. Konty (2009). A scan statistic for continuous data based on the normal probability model. *International Journal of Health Geographics* 8(1), 58.
- Loader, C. R. (1991). Large-deviation approximations to the distribution of scan statistics. *Advances in Applied Probability* 23(4), 751–771.
- Naus, J. I. (1982). Approximations for distributions of scan statistics. *J. Amer. Statist. Assoc.* 77(377), 177–183.
- Perone Pacifico, M., C. Genovese, I. Verdinelli, and L. Wasserman (2004). False discovery control for random fields. *J. Amer. Statist. Assoc.* 99(468), 1002–1014.
- Perone Pacifico, M., C. Genovese, I. Verdinelli, and L. Wasserman (2007). Scan clustering: a false discovery approach. *J. Multivariate Anal.* 98(7), 1441–1469.
- Picard, F., J.-C. Cadoret, B. Audit, A. Arneodo, A. Alberti, C. Battail, L. Duret, and M.-N. Prioleau (2014). The spatiotemporal program of dna replication is associated with specific combinations of chromatin marks in human cells. *PLoS Genet* 10(5), e1004282.
- Pipa, G., S. Grün, and C. van Vreeswijk (2013). Impact of spike train autostructure on probability distribution of joint spike events. *Neural Computation* 25, 1123–1163.
- Pouzat, C. and C. Chaffiol (2009). Automatic spike train analysis and report generation. an implementation with r, r2html and star. *Journal of Neuroscience Methods* 181, 119–144.
- Rivera, C. and G. Walther (2013). Optimal detection of a jump in the intensity of a poisson process or in a density with likelihood ratio statistics. *Scandinavian Journal of Statistics* 40(4), 752–769.

- Robin, S., F. Rodolphe, and S. Schbath (2005). *DNA, words and models*. Cambridge University Press, Cambridge. Translated from the 2003 French original.
- Romano, J. P. and M. Wolf (2005a). Exact and approximate stepdown methods for multiple hypothesis testing. *J. Amer. Statist. Assoc.* 100(469), 94–108.
- Romano, J. P. and M. Wolf (2005b). Exact and approximate stepdown methods for multiple hypothesis testing. *J. Amer. Statist. Assoc.* 100(469), 94–108.
- Shimazaki, H. and S. Shinomoto (2010). Kernel bandwidth optimization in spike rate estimation. *Journal of Computational Neuroscience* 29, 171–182.
- Siegmund, D. O., N. R. Zhang, and B. Yakir (2011). False discovery rate for scanning statistics. *Biometrika* 98(4), 979–985.
- Tuleau-Malot, C., A. Rouis, F. Grammont, and P. Reynaud-Bouret (2014). Multiple tests based on a Gaussian approximation of the Unitary Events method. *Neural Computation* 26(7).
- Walther, G. (2010, 04). Optimal and fast detection of spatial clusters with scan statistics. *Ann. Statist.* 38(2), 1010–1033.
- Westfall, P. H. and S. S. Young (1993). *Resampling-Based Multiple Testing*. Wiley. Examples and Methods for P -Value Adjustment.
- Yekutieli, D. (2008). False discovery rate control for non-positively regression dependent test statistics. *J. Statist. Plann. Inference* 138(2), 405–415.

A Supplementary material

A.1 Step-down improvement

Let us first introduce some notation. For any fixed subset \mathcal{C} of \mathcal{X}_η , $F_{\theta_0, \mathfrak{N}}^{\mathcal{C}}$ is the conditional c.d.f. given \mathfrak{N} of $\inf_{x \in \mathcal{C}} p_\eta(x)$. Then we define

$$\mathcal{N}(\mathcal{C}) = \{x \in \mathcal{X}_\eta, F_{\theta_0, \mathfrak{N}}^{\mathcal{C}}(p_\eta(x)) \leq \alpha\}.$$

The theoretical step-down algorithm can then be derived as follows:

- Step 0 : compute $\mathcal{R}^0 = \emptyset$ given by (8).
- Step $j \geq 1$: compute $\mathcal{R}^j = \mathcal{N}((\mathcal{R}^{j-1})^c)$

The step-down rejection set $\mathcal{R}_\eta^{\text{step-down}}$ is then defined as the limit of \mathcal{R}^j when j tends to infinity. Note that $\mathcal{R}^1 = \mathcal{R}_\eta^m$ is in fact given by (8) and that the sequence of \mathcal{R}^j is increasing.

The fact that this algorithm ends almost surely in a finite number of steps is just an easy consequence of the fact that for a given realization, there is only a finite number of possible values for the p -value process, which is in fact upper bounded by the size of the partition τ . Indeed, at each round, \mathcal{R}^j cannot decrease and if it increases, it is by absorbing sets of the type $\{x \in \mathcal{X}_\eta, p_\eta(x) = a\}$ for some a .

One can prove that such an algorithm guarantees a controlled FWER following the ideas of (Goeman and Solari, 2010). More precisely if one defines the event

$$\mathcal{E} = \{J_0^\eta \cap \mathcal{N}(J_0^\eta) = \emptyset\},$$

one has by (\mathcal{P}) (see Section 2.6) that

$$\begin{aligned} P_{\theta, \lambda}(\mathcal{E}^c) &= P_{\theta, \lambda}(\exists x \in J_0^\eta, F_{\theta_0, \mathfrak{N}}^{J_0^\eta}(p_\eta(x)) \leq \alpha) \\ &= P_{\theta_0, \lambda}(\exists x \in J_0^\eta, F_{\theta_0, \mathfrak{N}}^{J_0^\eta}(p_\eta(x)) \leq \alpha) \\ &= E_\lambda \left[P_{\theta_0} \left(F_{\theta_0, \mathfrak{N}}^{J_0^\eta}(\inf_{x \in J_0^\eta} p_\eta(x)) \leq \alpha \mid \mathfrak{N} \right) \right] \leq \alpha. \end{aligned}$$

The last point comes from the fact that $F_{\theta_0, \mathfrak{N}}^{J_0^\eta}(\inf_{x \in J_0^\eta} p_\eta(x))$ can be seen as the p -value of the test based on the statistics $-\inf_{x \in J_0^\eta} p_\eta(x)$ and described for instance in Lemma 1 of (Fromont et al., 2016), which implies that (4) holds even for the p -value $F_{\theta_0, \mathfrak{N}}^{J_0^\eta}(\inf_{x \in J_0^\eta} p_\eta(x))$. This is the single-step property of (Goeman and Solari, 2010).

For the monotony property of (Goeman and Solari, 2010), remark that if $\mathcal{C} \subset \mathcal{C}'$ then $F_{\theta_0, \mathfrak{N}}^{\mathcal{C}} \leq F_{\theta_0, \mathfrak{N}}^{\mathcal{C}'}$ and $\mathcal{N}(\mathcal{C}') \subset \mathcal{N}(\mathcal{C})$. Hence for all j , if $J_0^\eta \subset (\mathcal{R}^j)^c$, $J_0^{\eta+1} = \mathcal{N}([\mathcal{R}^j]^c) \subset \mathcal{N}(J_0^\eta)$.

Now, on \mathcal{E} , $\mathcal{N}(J_0^\eta) \cap J_0^\eta = \emptyset$ and therefore if $J_0^\eta \subset (\mathcal{R}^j)^c$, $J_0^\eta \subset (\mathcal{R}^{j+1})^c$ on the same event. Since $J_0^\eta \subset (\mathcal{R}^0)^c = \mathcal{X}_\eta$, one has with probability larger than $1 - \alpha$ that $J_0^\eta \subset \lim_j \mathcal{R}^j = \mathcal{R}_\eta^{\text{step-down}}$, which means that $\text{FWER}_{\theta, \lambda}^\eta(\mathcal{R}_\eta^{\text{step-down}}) \leq \alpha$.

A.2 Other kernel-based single test statistics

A.2.1 Homogeneity test statistics

Two sided test statistics. The starting point is that, conditionally on the event $\{N([0, 1]) = n\}$, the points in N form an n i.i.d. sample with density $f = 1 + \theta$. Hence, for any translation invariant kernel $K_h : [0, 1] \rightarrow \mathbb{R}$ of bandwidth h , we have

$$\widehat{f}_h(s) = \frac{1}{n} \sum_{T \in N \cap I_\eta(x)} K_h(s - T),$$

is an unbiased consistent estimate (when n tends to infinity) of

$$f_h(s) = \int_{I_\eta(x)} K_h(s - t) f(t) dt.$$

Consequently, the corresponding U -statistics given by

$$\frac{1}{n(n-1)} \sum_{\substack{T' \in N \cap I_\eta(x) \\ T' \neq T}} K_h(T - T'), \quad (11)$$

is an unbiased estimate of:

$$\int \int_{I_\eta(x)^2} K_h(s - t) f(t) f(s) ds \xrightarrow{h \rightarrow 0} \|f\|_{I_\eta(x)}^2,$$

under standard conditions on the kernel K_h . Actually, in a testing framework, we do not want to estimate $\|f\|_{I_\eta(x)}^2$ *stricto sensu*: here, we do not need to consider small values for h , we even recommend to take $h = \eta$ when it is unknown. This choice would not be appropriate for estimation, but seems sufficient for testing. Moreover if $\mathcal{H}_0 \{I_\eta(x)\}$ holds, $f = 1$ on $I_\eta(x)$ and $\|f\|_{I_\eta(x)}^2 = \eta$. So we propose the following statistics

$$S_\eta(x) = \left| \frac{1}{n(n-1)} \sum_{\substack{T' \in N \cap I_\eta(x) \\ T' \neq T}} K_h(T - T') - \eta \right|,$$

and the local null hypothesis is rejected for high values of $S_\eta(x)$.

One-sided test statistic. In this case we focus on what is happening when $f > 1$. Hence the target quantity to estimate is

$$\int \int_{I_\eta(x)^2} f(s) \max(f(s), 1) ds.$$

Here we consider the following density estimator:

$$\widehat{f}_h(s) = \frac{1}{(n-1)} \sum_{\substack{T' \in N \cap I_\eta(x) \\ T' \neq s}} K_h(T' - s),$$

which leads to the following statistics:

$$S_\eta(x) = \frac{1}{n} \sum_{T \in N \cap I_\eta(x)} \max\left(\widehat{f}_h(T), 1\right).$$

If there exists a t in $I_\eta(x)$ such that $f(t) > 1$ and if f is smooth enough such that this holds on a neighborhood, $E[\max(f(T), 1)\mathbf{1}_{\{T \in I_\eta(x)\}}]$ should be larger than η , whereas its value should be smaller than η if there are no such t .

A.2.2 Two-sample test statistics (one-sided)

Applying a reasoning similar to above, the quantity

$$\frac{1}{N([0, 1]) - 1} \sum_{\substack{T' \in N \cap I_\eta(x) \\ T' \neq T}} K_h(T - T')\varepsilon_{T'},$$

is a good estimate of $\theta(T)\lambda(T)$, for each $T \in N$. If one wants to reject only when there exists a $t \in I_\eta(x)$ such that $\theta(t) > 0$, then one needs to take the positive part of the previous quantity and integrate it with respect to $T \in N \cap I_\eta(x)$. We do not want to multiply again by $\varepsilon(T)$ because we do not want to multiply by θ twice. This leads to

$$S_\eta(x) = \frac{1}{N([0, 1])} \sum_{T \in N \cap I_\eta(x)} \max\left(\frac{1}{N([0, 1]) - 1} \sum_{\substack{T' \in N \cap I_\eta(x) \\ T' \neq T}} K_h(T - T')\varepsilon_{T'}, 0\right),$$

which estimates $E_\lambda[\max(\theta(T), 0)\mathbf{1}_{\{T \in I_\eta(x)\}}]$ for T of density proportional to λ .

B Proofs

B.1 Proof of Theorem 3.1

Computations and arguments similar to Section A.1 gives that

$$\begin{aligned}
\text{FWER}_{\theta,\lambda}^{\eta}(\mathcal{R}_{\eta}^m) &= E_{\lambda} [P_{\theta} (\exists x \in J_0^{\eta}, q_{\eta}(x) \leq \alpha \mid \mathfrak{N})] \\
&= E_{\lambda} \left[P_{\theta} \left(F_{\mathfrak{N}}^m \left(\inf_{x \in J_0^{\eta}} p_{\eta}(x) \right) \leq \alpha \mid \mathfrak{N} \right) \right] \\
&= E_{\lambda} \left[P_{\theta_0} \left(F_{\mathfrak{N}}^m \left(\inf_{x \in J_0^{\eta}} p_{\eta}(x) \right) \leq \alpha \mid \mathfrak{N} \right) \right] \\
&\leq E_{\lambda} \left[P_{\theta_0} \left(F_{\mathfrak{N}}^m \left(\inf_{x \in \mathcal{X}_{\eta}} p_{\eta}(x) \right) \leq \alpha \mid \mathfrak{N} \right) \right] \leq \alpha,
\end{aligned}$$

which concludes the proof.

B.2 Proof of Theorem 3.5

We follow the methodology introduced in Blanchard et al. (2014) by applying their Theorem 4.1 in the "Poisson case", as discussed in Example 2.5 and Section 4.2.4 therein. While the measurability conditions can be easily checked because the p -value process is càdlàg, the result will be proved if we show that the p -value process is finite-dimensional strong PRDS on J_0^{η} , as defined in Blanchard et al. (2014). For the homogeneity test, Lemma A.2 in Blanchard et al. (2014) shows that the p -value process is finite-dimensional strong PRDS on any subset, and thus also on J_0^{η} . For the two-sample test, this condition takes the following form: for any $q \geq 1$, for any $(x_j)_{1 \leq j \leq q} \in (\mathcal{X}_{\eta})^q$ such that $x_1 \in J_0^{\eta}$, for any nonincreasing set $D \subset \mathbb{N}^q$, for any $n \geq 0$,

$$\begin{aligned}
&P_{\theta,\lambda} (S_{\eta}(x_j)_{1 \leq j \leq q} \in D \mid S_{\eta}(x_1) = n + 1) \\
&\leq P_{\theta,\lambda} ((S_{\eta}(x_j)_{1 \leq j \leq q} \in D \mid S_{\eta}(x_1) = n). \tag{12}
\end{aligned}$$

Above, a nonincreasing set is defined as a subset $D \subset \mathbb{N}^q$ such that for all $(m, m') \in \mathbb{N}^q$ with $m_j \leq m'_j$ for $1 \leq j \leq q$, we have $m' \in D \Rightarrow m \in D$.

Now, to establish (12), we work *conditionally* on the joint Poisson process N . In that case, the random part of $S_{\eta}(x) = \sum_{T \in N \cap I_{\eta}(x)} (2\varepsilon_T - 1)$ is only carried by the set of marks $(\varepsilon_T)_{T \in N}$. Let $b_T = 2\varepsilon_T - 1$ for $T \in N$ and

$$D' = \left\{ (b_T)_{T \in N} \in \{0, 1\}^N : \left(\sum_{T \in N \cap I_{\eta}(x_j)} b_T \right)_{1 \leq j \leq q} \in D \right\}.$$

Since D is a nonincreasing set, so is D' and thus there exists some subset $\mathcal{T} \subset N$ such that $D' = \{(b_T)_{T \in N} \in \{0, 1\}^N : \forall T \in \mathcal{T}, b_T = 0\}$. Furthermore, since

$(b_T)_{T \in N \cap I_\eta(x_1)}$ is independent of $(b_T)_{T \in N \setminus I_\eta(x_1)}$, we have

$$\begin{aligned} & P_\theta \left((S_\eta(x_j))_{1 \leq j \leq q} \in D \mid S_\eta(x_1) = n, N \right) \\ &= P_\theta \left(\forall T \in \mathcal{T}, b_T = 0 \mid \sum_{T \in N \cap I_\eta(x_1)} b_T = n, N \right) \\ &= P_\theta (\forall T \in \mathcal{T} \setminus I_\eta(x_1), b_T = 0) \end{aligned} \quad (13)$$

$$\begin{aligned} & \times P_\theta \left(\forall T \in \mathcal{T} \cap I_\eta(x_1), b_T = 0 \mid \sum_{T \in N \cap I_\eta(x_1)} b_T = n, N \right) \\ &= P_\theta (\forall T \in \mathcal{T} \setminus I_\eta(x_1), b_T = 0) \times \binom{U-V}{n} / \binom{U}{n}, \end{aligned} \quad (14)$$

by denoting $U = \#N \cap I_\eta(x_1)$, $V = \#\mathcal{T} \cap I_\eta(x_1)$ and by an elementary combinatorial argument using that $b_T, T \in N \cap I_\eta(x_1)$, are i.i.d. $\mathcal{B}(1/2)$ variables. Finally, since $\binom{U-V}{n} / \binom{U}{n} = \binom{U-n}{V} / \binom{U}{V}$, the right-hand-side of (14) is nonincreasing w.r.t. n and the relation (12) is proved.

B.3 Proof of Theorem 5.1

Thanks to (\mathcal{P}) (see Section 2.6) that is used for the observed sample and the fact that the distribution of the resampled ε^b for $b = 1, \dots, B$, is fixed whatever θ , we have that

$$\begin{aligned} \text{FWER}_{\theta, \lambda}^\eta(\hat{\mathcal{R}}_\eta^m) &= E_\lambda \left[P_\theta \left(\exists x \in J_0^\eta, \hat{q}_\eta(x) \leq \alpha \mid \mathfrak{N} \right) \right] \\ &= E_\lambda \left[P_\theta \left(\frac{1}{B+1} \left(1 + \sum_{b=1}^B \mathbf{1}_{\{m^b \leq \inf_{x \in J_0^\eta} \hat{p}_\eta(x)\}} \right) \leq \alpha \mid \mathfrak{N} \right) \right] \\ &= E_\lambda \left[P_{\theta_0} \left(\frac{1}{B+1} \left(1 + \sum_{b=1}^B \mathbf{1}_{\{m^b \leq \inf_{x \in J_0^\eta} \hat{p}_\eta(x)\}} \right) \leq \alpha \mid \mathfrak{N} \right) \right] \\ &\leq E_\lambda \left[P_{\theta_0} \left(\frac{1}{B+1} \left(1 + \sum_{b=1}^B \mathbf{1}_{\{m^b \leq \inf_{x \in \mathcal{X}_\eta} \hat{p}_\eta(x)\}} \right) \leq \alpha \mid \mathfrak{N} \right) \right]. \end{aligned}$$

As a consequence, by using the notation m^0 , we have

$$\text{FWER}_{\theta, \lambda}^\eta(\hat{\mathcal{R}}_\eta^m) \leq E_\lambda \left[P_{\theta_0} \left(\frac{1}{B+1} \left(1 + \sum_{b=1}^B \mathbf{1}_{\{m^b \leq m^0\}} \right) \leq \alpha \mid \mathfrak{N} \right) \right].$$

Since, under P_{θ_0} , the random variables $\varepsilon^0, \dots, \varepsilon^B$ are i.i.d. Rademacher, the process vector $([\hat{p}_\eta^b(x)]_{x \in \mathcal{X}_\eta})_{b=0, \dots, B}$ is exchangeable (under P_{θ_0} and cond. on \mathfrak{N}), which




Figure 9: Simulation setting for the homogeneity case. Signal function $\theta(\cdot)$ that equals 0 under the null hypothesis, and θ^* elsewhere.

in turn implies that the vector $(m^b)_{b=0,\dots,B}$ is exchangeable (under P_{θ_0} and cond. on \mathfrak{N}). Therefore, applying Lemma 1 of (Romano and Wolf, 2005b), we obtain

$$P_{\theta_0} \left(\frac{1}{B+1} \left(1 + \sum_{b=1}^B \mathbf{1}_{\{m^b \leq m^0\}} \right) \leq \alpha \mid \mathfrak{N} \right) \leq \alpha,$$

and it remains to integrate in \mathfrak{N} to conclude.