



A note on online change point detection

Yi Yu^a, Oscar Hernan Madrid Padilla^b, Daren Wang^c, and Alessandro Rinaldo^d

^aDepartment of Statistics, University of Warwick, Coventry, UK; ^bDepartment of Statistics, University of California, Los Angeles, Los Angeles, California, USA; ^cDepartment of ACMS, University of Norte Dame, Notre Dame, Indiana, USA; ^dDepartment of Statistics & Data Science, University of Texas at Austin, Austin, Texas, USA

ABSTRACT

We investigate sequential change point estimation and detection in univariate nonparametric settings, where a stream of independent observations from sub-Gaussian distributions with a common variance factor and piecewise-constant but otherwise unknown means are collected. We develop a simple CUSUM-based methodology that provably control the probability of false alarms or the average run length while minimizing, in a minimax sense, the detection delay. We allow for all the model parameters to vary in order to capture a broad range of levels of statistical hardness for the problem at hand. We further show how our methodology is applicable to the case in which multiple change points are to be estimated sequentially.

ARTICLE HISTORY

Received 21 October 2022

Revised 7 September 2023

Accepted 17 October 2023

KEYWORDS

Online change point detection; cumulative sum statistics

1. INTRODUCTION

We investigate what may arguably be regarded as one of the most basic nonparametric online change point settings, whereby a sequence of independent, univariate sub-Gaussian random variables with known variance factor and unknown piecewise-constant means are observed sequentially. We seek to determine in an online manner, that is, each time we acquire a new observation, whether the data collected so far provide sufficient evidence to conclude that the mean of the distribution has changed at the present time or in the near past. The quality of any online procedure deployed for such a task is characterized based on the type I error, that is, the probability of incorrectly declaring that a change has taken place, and on the delay it incurs before correctly identifying a change point. Ideally, a good online procedure should guarantee a small false alarm probability while suffering only a minimal detection delay. It should be intuitively clear that these two features are at odds with each other: A procedure with a small probability of false alarm is likely to react slowly to even relatively big changes in the mean of the underlying distribution, thus producing a large detection delay. Vice versa, a methodology that is very sensitive to fluctuations in the data is unlikely to reliably discriminate between noise and signal and is thus prone to raising false alarms. In order to characterize this tradeoff, we keep track explicitly of all the parameters affecting the

CONTACT Y. Yu yi.yu.2@warwick.ac.uk Department of Statistics, University of Warwick, Coventry, UK

© 2023 The Author(s). Published with license by Taylor & Francis Group, LLC

This is an Open Access article distributed under the terms of the Creative Commons Attribution License (<http://creativecommons.org/licenses/by/4.0/>), which permits unrestricted use, distribution, and reproduction in any medium, provided the original work is properly cited. The terms on which this article has been published allow the posting of the Accepted Manuscript in a repository by the author(s) or with their consent.

difficulty of the change point detection task, such as the sub-Gaussian variance factor, the magnitude and time of the distributional change and the targeted type I error. Our goal is to develop an online procedure that provably works under the most unfavorable settings for which inference is just barely possible.

We begin by formalizing the problem with a general assumption used throughout.

Assumption 1.1. Assume that $\{X_1, X_2, \dots\}$ is a sequence of independent random variables with unknown means $\mathbb{E}(X_i) = f_i, i = 1, 2, \dots$ and such that $\sup_{i=1,2,\dots} \|X_i\|_{\psi_2} \leq \sigma$.

We recall that for a random variable X , its Orlicz- ψ_2 -norm is defined as

$$\|X\|_{\psi_2} = \inf\{t > 0 : \mathbb{E}\{\exp(X^2/t^2)\} \leq 2\}.$$

At first, we consider the change point model for which the means of the observations change after collecting Δ —an unknown number of observations, by an unknown amount κ , as described in the next assumption. Extension to models allowing for multiple change points will be discussed in [Section 3](#).

Assumption 1.2. Assume that there exists a positive integer $\Delta \geq 1$ such that

$$f_1 = \dots = f_\Delta \neq f_{\Delta+1} = f_{\Delta+2} = \dots.$$

In addition, let

$$\kappa = |f_\Delta - f_{\Delta+1}|.$$

We will write the probability of any event with respect to any distribution consistent with [Assumptions 1.1](#) and [1.2](#) as $\mathbb{P}_\Delta\{\cdot\}$ and, similarly, we will use $\mathbb{E}_\Delta\{\cdot\}$ for the corresponding expectation. With a slight abuse of notation, we describe the case in which the observations have constant means by setting $\Delta = \infty$ and will use the corresponding notation $\mathbb{P}_\infty\{\cdot\}$ and $\mathbb{E}_\infty\{\cdot\}$. With this notation, we remark that the change point is $\Delta + 1$ and Δ is the sample size of the observations from the pre-change distribution.

Our main goal is to design an online procedure that is provably able to detect a change point soon after time Δ and with a controlled false alarm probability, denoted throughout as α . In detail, an online change point detection procedure is an extended stopping time $\hat{t}$ taking values in $\mathbb{N} \cup \{\infty\}$ with respect to the natural filtration generated by the data. The false alarm probability of an online change point procedure is given as

$$\mathbb{P}_\Delta\{\hat{t} \leq \Delta\} \text{ for any } \Delta < \infty \quad \text{and} \quad \mathbb{P}_\infty\{\hat{t} < \infty\} \text{ otherwise,}$$

and the detection delay is the random variable

$$(\hat{t} - \Delta)_+,$$

which is only defined provided that $\Delta < \infty$. We will develop procedures that guarantee that (i) $\mathbb{P}_\infty\{\hat{t} < \infty\} \leq \alpha$, for a user-defined target false alarm probability α or $\mathbb{E}_\infty(\hat{t}) \geq \gamma$ for a user-defined average run length value γ , and (ii) that, at the same time, for all $\Delta < \infty$, $(\hat{t} - \Delta)_+$ is minimal.

The setting described in [Assumptions 1.1](#) and [1.2](#) allows one to completely characterize the hardness of the problem—measured both by the false alarm probability and the detection delay—as a function of the upper bound σ on the fluctuations, the target probability of false alarm α , the pre-change sample size Δ and the jump size κ .

Intuitively, the difficulty of the change point detection task is increasing in σ and decreasing in Δ , α and κ . To formally capture this important aspect of the problem, we implicitly assume a sequence of change point models with respect to which the data may have originated. Accordingly, the parameters defining the statistical task at hand, namely the quadruple $(\Delta, \alpha, \sigma^2, \kappa)$, are not fixed but should be instead viewed as sequences, expressing a spectrum of levels of difficulty of the problem we are interested in.

We make the following contributions:

- We develop a CUSUM-based procedure with a false alarm probability control, yielding a detection delay of order

$$\sigma^2 \kappa^{-2} \log(\Delta/\alpha) \quad (1.1)$$

with probability at least $1 - \alpha$, for all $\Delta < \infty$ satisfying the signal-to-noise ratio condition

$$\Delta \kappa^2 \sigma^{-2} \geq \log(\Delta/\alpha). \quad (1.2)$$

Interestingly, the above expression matches exactly the signal-to-noise ratio quantity for the offline version of the same change point detection problem, as demonstrated for example in Wang, Yu, and Rinaldo (2020). We elaborate further on this connection later on in Section 4.2.

- We show that a straightforward modification of our procedure guarantees that, for any target average run length $\gamma \geq \Delta$, $\mathbb{E}_\infty(\hat{t}) \geq \gamma$ and, at the same time, with probability $1 - \gamma^{-1}$ the detection delay is of order $\sigma^2 \kappa^{-2} \log(\gamma)$, for all $\Delta \leq \gamma$ such that $\Delta \kappa^2 \sigma^{-2} \geq \log(\gamma)$.
- We construct lower bounds indicating that the magnitude of the detection delay (1.1) is minimax rate-optimal whenever the signal-to-noise ratio condition (1.2) is in effect.
- We generalize our procedure to the case of multiple change points and show that, in this setting, the signal-to-noise condition (1.2) is in fact necessary for online change point localization.
- We discuss variants of our methodology that incur smaller computational costs. Chen, Wang, and Samworth (2020) mentioned that the computational cost of an online procedure should be of linear order of the number of time points. We would like to emphasize that this claim holds when both the before and after change point distributions are exactly known, which is essentially the situation discussed in Chen, Atev, and Lerman (2009). On the contrary, in our paper, we deal with sub-Gaussian distributions with unknown means. In this situation, we are not aware of nor expect to see any theoretically justified methods with linear order computational costs.

The paper is organized as follows. Section 1.1 provides a review of related work on online and offline change point detection. The main algorithms are then presented in Section 2. In particular, Section 2.1 contains the theoretical guarantees of the main algorithm, Section 2.2 discusses the variants, and some practical implementation aspects are investigated in Section 2.3. Extensions to multiple change points settings are the subject

of Section 3. Section 4 then provides a lower bound on the detection delay which shows that our algorithm is nearly optimal. The paper concludes with a discussion in Section 6. All the proofs are deferred to the Appendices.

1.1. Relevant literature

Wald (1945), as a prelude of the sequential analysis, kicked off the statistical research on online change point detection problems. A famous extension of Wald (1945) is the CUSUM statistic proposed in Page (1954). The optimality of Wald (1945) and Page (1954) was, to the best of our knowledge, studied first in Lorden (1971), which showed that among all the estimators which have average run lengths lower bounded by γ , the optimal detection delay rate is of order $\log(\gamma)/\text{KL}(F_0, F_1)$, as $\gamma \rightarrow \infty$, where $\text{KL}(\cdot, \cdot)$ is the Kullback–Leiber divergence. Moustakides (1986) and Ritov (1990) reiterated this minimax result and showed that in the optimality framework studied in Lorden (1971), the CUSUM statistic is optimal. Similar results have also been derived in the same framework using a change-of-measure argument under more general assumptions in Lai (1981), Lai (1998), and Lai (2001), among others. In almost all of the second half of the 20th century, the research on online change point detection focused on optimizing the expected delay time. The motivations back then were mainly from the manufactory sector, with quality control as the center applications. The data studied were mainly univariate sequences, and the results were almost all asymptotic. We refer readers to Lai (1995, 2001) for comprehensive reviews.

Before proceeding, we would like to emphasize that there is a fundamental difference between the optimality results derived in the aforementioned work and the ones developed by us in Section 4. In short, when considering the minimax lower bounds, the previous work only allows the change point location to vary and the optimality is derived in an asymptotic sense, by letting the lower bound of the average run length diverge. In this paper, we let all model parameters vary with the location of the change point when deriving the minimax lower bounds and allow for fixed sample arguments.

The second act of online change point research was kicked off by Chu, Stinchcombe, and White (1996), who formally stated the existence of “noncontamination” data that one has a training data set of size m , that is, $X_i \sim F_0, i = 1, \dots, m$. The theoretical results built upon the above assumption are asymptotic in the sense that one lets m grow unbounded. One can control the type I error with this noncontamination condition. Since Chu, Stinchcombe, and White (1996), a large number of papers have been produced in this line of work, including univariate mean change (e.g., Aue and Horváth 2004; Kirch 2008), linear regression coefficients change (e.g., Aue, Horváth, and Reimherr 2009; Husková and Kirch 2012), multivariate mean and/or variance change (e.g., Mei 2010), univariate nonparametric change (e.g., Husková, Kirch, and Meintanis 2010; Hlávka et al. 2016; Desobry, Davy, and Doncarli 2005), and Bayesian online change point detection (e.g., Fearnhead and Liu 2007), to name but a few. More recent work includes He et al. (2018), which studied the sequential change point detection in a sequence of random graphs. Kirch and Weber (2018) used estimating equations as a unified framework to include the location shift, linear regression, and autoregressive online change point detection. Kurt, Yilmaz, and Wang (2018) converted different high-

dimensional and/or nonparametric data to a univariate statistic and used geometric entropy minimization methods to define an acceptance region. Chen (2019) constructed similarity measures via K -nearest neighbor estimators and then proposed a counting-based statistic to conduct sequential change point detection. Dette and Gösmann (2020) proposed a general framework for sequential change point detection case and obtained a limiting distribution of the proposed statistics. The framework can be used to handle high-dimensional and nonparametric cases. Gösmann, Kley, and Dette (2019) exploited a likelihood ratio based method and shared similar core techniques with Dette and Gösmann (2020). Keshavarz, Michailidis, and Atchade (2018) considered online change point detection in a sequence of Gaussian graph models and obtained asymptotic type I and II error controls. Chen, Wang, and Samworth (2020) considered online change point detection in a sequence of Gaussian random vectors where the mean changes over time. Comprehensive monographs and survey papers include Siegmund (2013), Tartakovsky, Nikiforov, and Basseville (2014), and Namono et al. (2019).

We would like to mention Maillard (2019). This paper is the most relevant paper to ours, to the best of our knowledge, and has heavily inspired our paper. Maillard (2019) studied a univariate online mean change point detection problem and deployed the Laplace transform to control the probabilities of the events, on which the fluctuations are contained within desirable ranges, although the arguments thereof remain doubtful. Given the large probability events established based on the Laplace transforms, type I error controls and large probability detection delays were studied. A claim on the phase transition and robust analysis under multiple change points scenario were also available in Maillard (2019). There are a number of differences between this paper and Maillard (2019): (i) Instead of using the Laplace transform to establish large probability events, we summon the concentration inequalities for sub-Gaussian random variables, the union bound results, and the peeling arguments. It was pointed out in the Discussion in Maillard (2019), other more advanced tools including the peeling arguments may improve the results by changing logarithmic terms to iterative logarithmic terms. We are, however, skeptical about the feasibility of such claim. (ii) In addition to the type I errors, which are available in Maillard (2019), we also provide average run length results and a parallel set of results by setting a lower bound for the average run lengths. This is a common practice in applications and is widely used in the existing literature (e.g., Lai 1981). (iii) Despite that Maillard (2019) has much more modern arguments than papers in the 20th century, the results are presented in a more restrictive way. For instance, the “phase transition” and “detectability” are presented as a property of the location of the change point only. In this paper, we will exploit a signal-to-noise ratio, which is a function of the jump size, the variance and the change point location jointly. This setup enables further studies of high-dimensional data problems.

An important aspect of our problem that sets it aside from many of the contributions in the area is the fact that our procedure is not sensitive to the values of the pre- and post-change means, but only to their difference. Formally, we consider composite null and alternative hypotheses for the pre- and post-change distributions, respectively. This feature, combined with the nonparametric nature of our sub-Gaussianity assumption, which is not enough to specify likelihood functions, poses additional and significant challenges compared to the simpler case in which the pre- and post-change distributions

are known. From a statistical standpoint, it prevents us from relying on likelihood-ratio-type procedures, which are analytically simpler and have optimality properties. Computationally, it requires us to reevaluate our test statistic over past observations, adding to the computational burden and thus making the design of sequential procedures challenging.

2. DETECTION DELAY AND TYPE I ERROR CONTROLS

We are going to consider a simple online procedure, described in [Algorithm 1](#), that keeps track of the running maximum of the standard CUSUM statistics, formally defined next, and declares that a change point has occurred as soon as its value exceeds a pre-specified, time-dependent threshold. Later in the paper, we will discuss different variants of this simple method, but they all share this basic core principle.

Definition 1 (CUSUM statistic). Given a sequence $\{X_t\}_{t=1,2,\dots} \subset \mathbb{R}$ and a pair of integers $1 \leq s < t$, we define the CUSUM statistic and its population version as

$$\widehat{D}_{s,t} = \left| \left(\frac{t-s}{ts} \right)^{1/2} \sum_{l=1}^s X_l - \left\{ \frac{s}{t(t-s)} \right\}^{1/2} \sum_{l=s+1}^t X_l \right|$$

and

$$D_{s,t} = \left| \left(\frac{t-s}{ts} \right)^{1/2} \sum_{l=1}^s f_l - \left\{ \frac{s}{t(t-s)} \right\}^{1/2} \sum_{l=s+1}^t f_l \right|,$$

respectively.

Algorithm 1. Online change point detection vis CUSUM statistics

INPUT: $\{X_u\}_{u=1,2,\dots} \subset \mathbb{R}, \{b_u, u = 2, 3, \dots\} \subset \mathbb{R}$.

$t \leftarrow 1, \text{FLAG} \leftarrow 0$

while FLAG = 0 **do**

$t \leftarrow t + 1;$

FLAG = $1 - \prod_{s=1}^{t-1} 1\{\widehat{D}_{s,t} \leq b_t\};$

end while

OUTPUT: t .

[Algorithm 1](#) scans through the data sequence using the CUSUM statistic and a sequence of prespecified positive threshold values. For any time point $t \geq 2$, as long as there exists an integer $s \in [1, t)$, such that the corresponding CUSUM statistic $\widehat{D}_{s,t}$ exceeds the threshold b_t , we decide that a change in mean has occurred prior to the current time point t . [Algorithm 1](#) is written in the way that it will never terminate if there is no change point declared. In practice, [Algorithm 1](#) is terminated based on a stopping rule decided by the user and our theory accommodates this.

In the rest of this section, we will first focus on one version of the detection procedure, providing its detection delay and type I error control in [Section 2.1](#). We will then discuss two other common alternative procedures, together with their performances and

connections among all three procedures in [Section 2.2](#). To conclude this section, we will discuss some practical issues in [Section 2.3](#).

2.1. Analysis of Algorithm 1

In this section, we analyze the performance of [Algorithm 1](#). Here and in the rest of the paper, we adopt a high-dimensional framework of reference, whereby we implicitly assume a sequence of change point models with respect to which the data may have originated. Accordingly, the parameters defining the statistical task at hand, namely the pre-change sample size Δ when it is finite, the variance parameter σ , the magnitude κ of the change and the false alarm probability α are not fixed but are instead allowed to vary, in order to express a spectrum of levels of difficulty of the problem we are interested in. This is of course a convenient mathematical formalization and is not intended to represent any real-life situation. Since our theoretical guarantees will be formulated based on finite sample bounds, this formalism does not pose any issue.

We begin by specifying the signal-to-noise ratio condition that we will require in our analysis.

Assumption 2.1. There exists a sufficiently large absolute constant $C_{\text{SNR}} > 0$ such that

$$\Delta \kappa^2 \sigma^{-2} \geq C_{\text{SNR}} \log(\Delta/\alpha).$$

It is worth emphasizing that the above condition should be interpreted as a constraint that is expected to be satisfied for all values of the parameters $(\Delta, \sigma^2, \kappa, \alpha)$ simultaneously. For example, assume the type I error level α and the jump size parameter κ to be fixed. Then, the above condition will constrain how large the variance parameter σ^2 is allowed to be in relation to Δ and, vice versa, how small Δ can be with respect to σ^2 . We will refer to any relationship that holds among $(\Delta, \sigma^2, \kappa, \alpha)$, such as the one expressed in [Assumption 2.1](#), as a scaling.

In our main result, we show that using the CUSUM statistic and under the signal-to-noise ratio condition in [Assumption 2.1](#), [Algorithm 1](#) is able to detect the change point with a detection delay of order up to $\sigma^2 \kappa^{-2} \log(\Delta/\alpha)$, with probability at least $1 - \alpha$, using time-varying thresholds of order $\sigma \log^{1/2}(t/\alpha)$. This rate is minimax optimal, as we show in [Section 4](#).

Theorem 2.2. Consider the settings described in [Assumption 1.1](#). Let $\alpha \in (0, 1)$ and $\hat{t}$ be the stopping time returned by [Algorithm 1](#) with inputs $\{X_t\}_{t=1,2,\dots}$ and $\{b_t\}_{t=2,3,\dots}$, where

$$b_t = 2^{3/2} \sigma \log^{1/2}(t/\alpha). \quad (2.1)$$

If $\Delta = \infty$, then

$$\mathbb{P}_\infty\{\hat{t} < \infty\} < \alpha. \quad (2.2)$$

Under [Assumption 1.2](#), we have

$$\mathbb{P}_\Delta\{\hat{t} \leq \Delta\} < \alpha, \quad (2.3)$$

for any $\Delta \geq 1$. If [Assumptions 1.2](#) and [2.1](#) both hold, then

$$\mathbb{P}_\Delta \left\{ \Delta < \hat{t} \leq \Delta + C_d \frac{\sigma^2 \log(\Delta/\alpha)}{\kappa^2} \right\} \geq 1 - \alpha, \quad (2.4)$$

where $C_d > 0$ is an absolute constant with $C_d < C_{\text{SNR}}$.

Remark 1. In relation to the procedure described in Lorden (1971), a direct consequence of combining Theorem 2.2 above and Theorem 2 in Lorden (1971) is the following. For $k \in \{1, 2, \dots\}$, let $\hat{t}_k$ be the output of applying Algorithm 1 to $\{X_k, X_{k+1}, \dots\}$ and let $\hat{t}_* = \min\{\hat{t}_k + k - 1, k = 1, 2, \dots\}$. We have that the average run length of $\hat{t}_*$ satisfies $\mathbb{E}_0(\hat{t}_*) \geq 1/\alpha$.

The proof of Theorem 2.2 is given in Section B. In our proofs, we have not optimized the constants, which we conjecture can be made smaller through a more refined analysis. As for the tuning parameter b_t used by the procedure, we have made the multiplicative constant $2^{3/2}$ explicit, though any number larger than 2 can be used: See (18). In practice, since usually the variance parameter σ is unknown, we recommend calibrating the values of b_t via simulations, assuming the availability of data from a pre-change distribution, as it is often the case in applications.

The guarantee (2.2) implies that, in the absence of any change points, the procedure will continue indefinitely with probability at least $1 - \alpha$. On the other hand, when a change point does exist, as assumed in Assumption 1.2, the bound (2.3) guarantees that the false alarm probability of our procedure is no larger than α , regardless of the actual value of the change point. Additionally, under the scaling on the signal-to-noise ratio given in Assumption 2.1, (2.4) provides a high-probability bound on the detection delay of order $\sigma^2 \kappa^{-2} \log(\Delta/\alpha)$, which is increasing in σ and Δ and decreasing in α and κ . In particular, combining (2.4) with Assumption 2.1, we can see that the smaller α —the tolerance on the type I error—is, the larger the required signal-to-noise ratio and the detection delay are. As suggested by the minimax lower bound given below in Proposition 4.1, our detection delay bound appears to be optimal. We refer readers to Section 4 for a detailed discussion of the optimality of our results and their relation to the existing literature. Finally, we remark the expression of detection delay bound is nearly identical to the one of the localization error found by Wang, Yu, and Rinaldo (2020) for the problem of estimating the change points in offline settings, which too is nearly minimax rate-optimal. While this result may not be surprising, it is far from obvious.

Theorem 2.2 could be equivalently stated using stopping times leading to a more practical interpretation. Such formulation is better suited to accommodate real-life situations, where the experiment cannot continue indefinitely and is instead ultimately terminated based on a predefined, possibly random, stopping rule. Specifically, the proof of Theorem 2.2 also implies that, with probability at least $1 - \alpha$ and uniformly over all finite random times T , it holds that

- $\hat{t} > T$ if $\Delta = \infty$ or if $T \leq \Delta$,
- $\Delta < \hat{t} \leq \Delta + C_d \frac{\sigma^2 \log(\Delta/\alpha)}{\kappa^2}$, if $T > \Delta + C_d \frac{\sigma^2 \log(\Delta/\alpha)}{\kappa^2}$.

Whenever T happens to be between Δ and $\Delta + C_d \sigma^2 \kappa^{-2} \log(\Delta/\alpha)$, then the procedure is only guaranteed to not raise a false alarm with probability at least $1 - \alpha$, at any time prior to and including Δ but may not be able to reach the correct decision that a change point has occurred since there are not enough observations between Δ and T to make such a determination.

The proof of [Theorem 2.2](#), as well as of most of the results in the paper, relies in a fundamental way on an auxiliary result given in [Lemma A.1](#), which shows that, with probability at least $1 - \alpha$, the maximal CUSUM statistic process $t \in \{2, 3, \dots\} \mapsto \max_{1 \leq s \leq t} \widehat{D}_{s,t}$ does not exceed the curved boundary of the form $t \mapsto 2^{3/2} \sigma \sqrt{\log(t)/\alpha}$. The proof also shows that, for any fixed $t \geq 2$, the order of $\max_{1 \leq s \leq t} \widehat{D}_{s,t}$ is no larger than $\sigma \sqrt{\log \log(t/\alpha)}$ with probability at least $1 - \alpha$. Interestingly, the iterated logarithmic scaling does not appear to apply to the maximal CUSUM process, but only to the CUSUM maximum at any fixed time t .

[Theorem 2.2](#) shows that [Algorithm 1](#) can detect the change point with probability at least $1 - \alpha$, provided that

$$\Delta \kappa^2 \sigma^{-2} \geq \log(\Delta/\alpha),$$

as required in [Assumption 1.2](#). In fact, this condition is essentially necessary for our procedure. More precisely, in our next result we show that if

$$\Delta \kappa^2 \sigma^{-2} \leq \log(\Delta/\alpha), \quad (2.5)$$

then, with probability at least $1 - \alpha$, [Algorithm 1](#) will never terminate and therefore will be unable to reliably detect a change point. That is, under the scaling (2.5), our procedure only delivers type I error control but is virtually powerless. We conjecture that the scaling in [Assumption 1.2](#) is required by *any* online procedure in order to guarantee a false alarm probability smaller than α and at the same time, non-trivial power (which in our problem would correspond to a probability no smaller than $1 - \alpha$ of terminating in a finite time when there exists a change point).

Proposition 2.3. *Fix an $\alpha \in (0, 1)$ and let [Assumptions 1.1](#) and [1.2](#) hold. Assume in addition that for any $\alpha \in (0, 1)$, there exists a positive constant $c_{\text{SNR}} > 0$ such that*

$$\Delta \kappa^2 \sigma^{-2} \leq c_{\text{SNR}} \log(\Delta/\alpha). \quad (2.6)$$

Let $\widehat{t}$ be the stopping time returned by [Algorithm 1](#) with inputs $\{X_t\}_{t=1,2,\dots}$ and $\{b_t\}_{t=2,3,\dots}$, where

$$b_t = C_b \sigma \log^{1/2}(t/\alpha), \quad t \geq 2,$$

with $C_b > 0$ being sufficiently large. Then,

$$\mathbb{P}_\Delta(\widehat{t} < \infty) < \alpha.$$

2.2. Variations of Algorithm 1

The framework for online change point detection we have analyzed so far is based on controlling the false alarm probability α , and the magnitude of the detection delay of

our procedure depends on the choice of the target probability; see (2.4). An alternative approach, considered in the literature and frequently adopted in practice, is to impose the weaker requirement that, for a target value $\gamma > 0$, $\mathbb{E}_\infty(\hat{t}) \geq \gamma$. That is, the expected false detection time, or average run length, of an online procedure should be no smaller than γ when in fact there is no change point. In practical applications in which the pre-change point distribution is known or can be estimated from historical data, online procedures are often calibrated through simulations by choosing γ according to this criterion (e.g., Lorden 1971; Lai 2001).

Below we show how Algorithm 1 can be tuned to yield this type of false alarm control, as well as a high probability bound on the detection delay when there is in fact a change point, where the magnitude of the detection delay scales logarithmically with γ .

Proposition 2.4. *Consider the settings described in Assumption 1.1. Let $\gamma \geq 2$ and $\hat{t}$ be the stopping time returned by Algorithm 1 with inputs $\{X_t\}_{t=1,2,\dots}$ and $\{b_t\}_{t=2,3,\dots}$, where*

$$b_t = C_b \sigma \log^{1/2} \left\{ 2^{1/3} (\gamma + 1) \right\} \quad (2.7)$$

for an absolute constant $C_b > 0$. If $\Delta = \infty$, then,

$$\mathbb{E}_\infty(\hat{t}) \geq \gamma.$$

Under Assumption 1.2, if

$$\gamma \geq \Delta \quad \text{and} \quad \Delta \kappa^2 \sigma^{-2} \geq C_{\text{SNR}} \log(\gamma), \quad (2.8)$$

where $C_{\text{SNR}} > 0$ is an absolute constant, then

$$\mathbb{P}_\Delta \left\{ \Delta < \hat{t} \leq \Delta + C_d \frac{\sigma^2 \log(\gamma)}{\kappa^2} \right\} \geq 1 - \gamma^{-1},$$

where $C_d > 0$ is an absolute constant.

The proof of Proposition 2.4 is in Appendix B. Just like with Theorem 2.2, it is immediate to formulate a version of Proposition 2.4 that is based on stopping times. For brevity, we refrain from providing details.

It is illustrative to compare the guarantees afforded by Theorem 2.2 and Proposition 2.4 when a change point exists. In this case, the signal-to-noise ratio condition defined in Assumption 2.1 is replaced by the analogous conditions shown in (2.8), where it is crucial that γ is an upper bound on Δ . In terms of the upper bounds on the detection delay, the difference is between $\log(\gamma)$ and $\log(\Delta/\alpha)$. When γ is of the same order as Δ , then these two detection delay upper bounds differ by a factor $\log(1/\alpha)$, which is of constant order in the fixed confidence settings where α is held constant. In terms of the values of the probability bounds, the difference is between γ^{-1} and α . This further suggests that, as long as

$$\gamma \asymp \alpha^{-1}, \quad \gamma \geq \Delta \quad \text{and} \quad \gamma \asymp \Delta,$$

these two strategies are equivalent in terms of controlling the detection delay. This connection has been studied before in a slightly different form, see e.g., Lai (1998).

In Proposition 2.4, if Assumption 1.2 holds but (2.8) does not, especially if $\gamma < \Delta$, then we can only reach the conclusion that $\mathbb{E}_\Delta(\hat{t}) \geq \gamma$, but no longer guarantee the

control on the type I error or on the detection delay. This is because the values of the threshold b_ν , defined in (2.7), is constant with respect to t and only depends on γ .

In another variant of Algorithm 1 that is often used in practice, researchers may only wish to control the false alarm probability over any interval of a prespecified length. This strategy may be preferred for computational reasons, as the computational cost of the procedure can be directly controlled by selecting an appropriate interval length, especially when training can easily be done with historical data. Below, we provide a parallel result to Theorem 2.2. To that effect, we require a new definition of the CUSUM statistic in Definition 2 and slightly modify Algorithm 1, resulting in Algorithm 2, whose computational cost is of order $O(t)$, when proceeding to time point t .

Definition 2. Given data $\{X_t\}_{t=1,2,\dots} \subset \mathbb{R}$ and a triplet of integers $1 \leq e < s < t$, we define the CUSUM statistic and its population counterpart as

$$\widehat{D}_{e,s,t} = \left| \left\{ \frac{t-s}{(s-e)(t-e)} \right\}^{1/2} \sum_{l=e+1}^s X_l - \left\{ \frac{s-e}{(t-s)(t-e)} \right\}^{1/2} \sum_{l=s+1}^t X_l \right|$$

and

$$D_{e,s,t} = \left| \left\{ \frac{t-s}{(s-e)(t-e)} \right\}^{1/2} \sum_{l=e+1}^s f_l - \sqrt{\frac{s-e}{(t-s)(t-e)}} \sum_{l=s+1}^t f_l \right|,$$

respectively.

Algorithm 2. Online change point detection 2.

INPUT: $\{X_u\}_{u=1,2,\dots} \subset \mathbb{R}, w \geq 2, b_w > 0$.

$t \leftarrow 1, \text{FLAG} \leftarrow 0$

while FLAG = 0 **do**

$t \leftarrow t + 1$;

FLAG = $1 - \prod_{s=\max\{t-w+1, 1\}}^{t-1} \mathbb{1} \left\{ \widehat{D}_{\max\{t-w, 0\}, s, t} \leq b_w \right\}$;

end while

OUTPUT: t .

Proposition 2.5. Consider the settings described in Assumption 1.1. Let $\alpha \in (0, 1)$ and $\widehat{t}$ be the stopping time returned by Algorithm 2 with inputs $\{X_t\}_{t=1,2,\dots}, w \geq 2$ and $b_w = 2^{1/2} \sigma \log^{1/2}(2w^2/\alpha)$. Then, when $\Delta = \infty$,

$$\sup_{v \geq 1} \mathbb{P}_\infty(v < \widehat{t} \leq v + w) < \alpha.$$

If Assumptions 1.2 and 2.1 hold and

$$w > C_w \sigma^2 \log(2w^2/\alpha) \kappa^{-2}, \quad (2.9)$$

with an absolute constant $C_w > 0$ then with an absolute constant $C_d > 0$,

$$\mathbb{P} \left\{ (\widehat{t} - \Delta)_+ \geq C_d \sigma^2 \log(2w^2/\alpha) \kappa^{-2} \right\} < \alpha.$$

The proof of Proposition 2.5 can be found in Appendix B. The strategy used in proving Proposition 2.5 can be seen as a mixture of the strategies used in Theorem 2.2 and

Proposition 2.4: Instead of controlling the type I error over the whole time course, it is enough to only control the type I error over intervals of length w . Comparing with [Proposition 2.4](#), the advantage here is that, if $w < \Delta$, we will be able to provide a high-probability bound on the detection delay. As for [Theorem 2.2](#), the advantage here is the same as that in [Proposition 2.4](#), that it might be handier in tuning parameter selection, while the price it pays is here that, when $w < \Delta$, then there is no guarantee that $\hat{t} > \Delta$, that is, preventing false alarms.

2.3. Practical and computational issues

In this section, we are to discuss two practical issues for the CUSUM statistics-based online procedures discussed above.

The first issue is the choice of tuning parameters. In [Algorithm 1](#), we need a sequence of tuning parameters $\{b_t\}$, whose theoretically justified value is $2^{3/2}\sigma \log^{1/2}(t/\alpha)$; see [\(2.1\)](#). The quantity α is specified by the user, but the sub-Gaussian parameter σ remains unknown. In some situations, the practitioner has access to independent copies of data generated from a model with no change points. Then, one may wish to set a limit of time, say T , and estimate the empirical type I errors over the time course $\{1, \dots, T\}$, in order to tune the thresholds. In this sense, the first variant we mentioned in [Section 2.2](#) is handier. The tuning parameter can be chosen by setting the average run length equal to a pre-specified γ .

The second practical issue is computational complexity. The CUSUM statistic given in [Definition 1](#) can be rewritten as

$$\hat{D}_{s,t} = \left\{ \frac{t}{s(t-s)} \right\}^{1/2} \left| \frac{s}{t} \sum_{i=1}^t X_i - \sum_{i=1}^s X_i \right|.$$

Using this equivalent expression, for each time t , one can store all partial sums $\{\sum_{i=1}^s X_i\}_{s=1}^t$ and the computational cost of [Algorithm 1](#) is therefore of order $O(t)$ but the storage is also of order $O(t)$. As an alternative, one may elect to recalculate the CUSUM statistic every time, in which case there is no requirement on storage but the computational cost increases to be of order $O(t^2)$.

To reduce the computational burden, a very simple strategy is to avoid calculating the maximal values of the CUSUM statistics at all integer pairs $1 \leq s < t$ and instead consider only pairs that are at a certain distance in time, say $h > 0$, a prespecified parameter picked by the user. The window width h can be regarded as the user's maximal tolerance on accuracy. More precisely, at each time $t \geq 2$, instead of maximizing the values of the CUSUM statistics over all integers $s \in [1, t)$, one could just calculate $\hat{D}_{t-h,t}$, for

$$h \asymp \frac{\sigma^2}{\kappa^2 \log(\Delta)} \quad \text{and} \quad t > h.$$

Then, in [Algorithm 1](#), it would only be enough to check if $\hat{D}_{t-h,t}$ exceeds an appropriate threshold. The computational complexity of this alternative is of order $O(t)$ and the storage cost is of order $O(1)$. The caveat of this alternative is that one needs to

carefully tune h , which is required to essentially have the same order of the magnitude as the detection delay. To address this issue, we propose to compute $\widehat{D}_{t-h,t}$ only over a geometrically increasing sequence of values for h , a tuning strategy that will only incur in an additional computational cost of order $O\{\log(t)\}$ when proceeding to time point t . The detailed procedure is given in [Algorithm 3](#). As it turns out, this strategy is not only computationally convenient, but it yields the same nearly optimal theoretical guarantees of [Theorem 2.2](#), as shown next.

Algorithm 3. Online change point detection—varying h .

INPUT: $\{X_u\}_{u=1,2,\dots} \subset \mathbb{R}, \{b_u, u = 1, 2, \dots\} \subset \mathbb{R}$.

$t \leftarrow 1, \text{FLAG} \leftarrow 0$

while FLAG = 0 **do**

$t \leftarrow t + 1;$

$J \leftarrow \lfloor \log(t) / \log(2) \rfloor;$

$j \leftarrow 0;$

while $j < J$ and FLAG = 0 **do**

$j \leftarrow j + 1;$

$s_j \leftarrow t - 2^{j-1};$

FLAG = $1\{\widehat{D}_{s_j,t} > b_t\};$

end while

end while

OUTPUT: t .

Corollary 2.6. Consider the settings described in [Assumption 1.1](#). Let $\alpha \in (0, 1)$ and $\widehat{t}$ be the stopping time returned by [Algorithm 3](#) with inputs $\{X_t\}_{t=1,2,\dots}$ and $\{b_t\}_{t=2,3,\dots}$, where

$$b_t = 2^{3/2} \sigma \log^{1/2}(t/\alpha), \quad t \geq 2.$$

If $\Delta = \infty$, then

$$\mathbb{P}_\infty\{\widehat{t} < \infty\} < \alpha.$$

Under [Assumption 1.2](#), we have

$$\mathbb{P}_\Delta\{\widehat{t} \leq \Delta\} < \alpha,$$

for any $\Delta \geq 1$. If [Assumptions 1.2](#) and [2.1](#) both hold, then

$$\mathbb{P}\left\{\Delta < \widehat{t} \leq \Delta + C_d \frac{\sigma^2 \log(\Delta/\alpha)}{\kappa^2}\right\} \geq 1 - \alpha,$$

for all $\Delta \geq 1$, where $C_d > 0$ is an absolute constant satisfying $C_d < C_{\text{SNR}}$.

3. MULTIPLE CHANGE POINTS

It is relatively straightforward to extend our methodology and analysis to deal with multiple change points. In this setting, the analyst still collects data sequentially and, with each new data point, makes a decision as to whether there exists evidence supporting the presence of a change point in the near past. But, unlike the single change point case

considered in the previous sections, the procedure is restarted each time a new change point is declared, until the experiment is terminated. In order to accommodate for the presence of multiple change points, we need a refined setup.

Assumption 3.1. Assume that there exists a collection of change points $1 = \eta_0 < \eta_1 < \eta_2 < \dots$, such that

$$f_{\eta_{k-1}} = \dots = f_{\eta_k-1} = \mu_k, \quad k \in \mathbb{N}^*,$$

where $\mu_k \neq \mu_{k+1}$ for each $k \in \mathbb{N}^*$. Let

$$\Delta = \inf_{k \in \mathbb{N}^*} (\eta_k - \eta_{k-1})$$

and, for each $k \in \mathbb{N}^*$,

$$\kappa_k = |\mu_k - \mu_{k+1}|.$$

Finally, set $\kappa = \inf_{k \in \mathbb{N}^*} \kappa_k$, assumed to be strictly positive.

In [Assumption 3.1](#), the minimal spacing between two consecutive change points Δ and the minimal jump size κ are natural generalizations of the analogous quantities defined in [Assumption 1.2](#) and indeed coincide with them when there is only one change point.

[Algorithm 4](#) is an immediate generalization of [Algorithm 1](#) and amounts to repeatedly applying [Algorithm 1](#) as soon as a new change point is found. Below we will demonstrate that, under the signal-to-noise ratio condition described in [Assumption 3.2](#), given a prespecified confidence level $\alpha \in (0, 1)$, the input parameters $\{b_u, u = 2, 3, \dots\}$ of the algorithm can be chosen so that, with probability at least $1 - \alpha$,

- the procedure will not declare any change point if in fact there is none;
- if there are change points, then, for any stopping time T with respect to the natural filtration induced by the data, the procedure will not return any change point if $T \leq \eta_1$ and will estimate the sequence $\{\eta_1, \dots, \eta_K\} = \{1, \dots, T\} \cap \{\eta_k\}_{k=1,2,\dots}$ of true change points prior to T . Furthermore, in the former case, the detection delays will be of order $O\left\{\sigma^2 \kappa_j^{-2} \log(\Delta/\alpha)\right\}, j = 1, \dots, K$, with high probability. It is worth noting that the number K will in general be random if the stopping time T is not deterministic.

Algorithm 4. Online change point detection—multiple change points.

INPUT: $\{X_u\}_{u=1,2,\dots} \subset \mathbb{R}, \{b_{e,u}, u = 2, 3, \dots, e = 0, 1, \dots, u-1\} \subset \mathbb{R}, \mathcal{C} = \emptyset$
 $e \leftarrow 0, t \leftarrow 1, \text{FLAG} \leftarrow 0$.

while there is a new data point **do**

$t \leftarrow t + 1$;

$\text{FLAG} = 1 - \prod_{s=1}^{t-1} \mathbb{1}\left\{\widehat{D}_{e,s,t} \leq b_{e,t}\right\}$;

if $\text{FLAG} = 1$ **then**

$\mathcal{C} \leftarrow \mathcal{C} \cup \{t\}$;

$\text{FLAG} \leftarrow 0$;

$e \leftarrow t$
end if
end while
OUTPUT: C .

The signal-to-noise condition that we required for the case of multiple change points is as follows.

Assumption 3.2. Assume that for any $\alpha \in (0, 1)$, there exists a sufficiently large absolute constant $C_{\text{SNR}} > 0$ such that

$$\Delta \kappa^2 \sigma^{-2} \geq C_{\text{SNR}} \log(\Delta/\alpha).$$

Assumptions 2.1 and 3.2 are identical in appearance, though, of course, the meaning of the quantities Δ and κ are different depending on the settings. In the main result of this section, we generalize the findings in (2.4) to the settings involving multiple change points, deriving essentially the same false alarm and detection delay bounds as in the only one change point case. In order to allow for a more transparent statement, we have expressed our results in terms of a uniform choice of stopping times for terminating the experiment.

Theorem 3.3. Consider the settings described in Assumption 1.1. Let $\alpha \in (0, 1)$ and C be the output of Algorithm 4 with inputs $\{X_t\}_{t=1,2,\dots}$ and $\{b_u, u = 2, 3, \dots\}$, where

$$b_u = 4\sigma \log^{1/2}(t/\alpha), \quad \forall u = 2, 3, \dots \quad (3.1)$$

The following holds with probability at least $1 - \alpha$, uniformly over all choices of finite stopping times T with respect to the natural filtration generated by $\{X_t\}_{t=1,2,\dots}$.

- If the random variables $\{X_t\}_{t=1,2,\dots}$ have constant means, then $C = \emptyset$.
- Under Assumptions 3.1 and 3.2, let $K = |\{1, \dots, T\} \cap \{\eta_k, k = 1, 2, \dots\}|$. If $K = 0$, then

$$C \cap \{1, \dots, T\} = \emptyset.$$

If $K \geq 1$, there exists an absolute constant $C_d > 0$ satisfying $C_d < C_{\text{SNR}}$ such that

$$C = \begin{cases} \{\hat{t}_1, \dots, \hat{t}_{K-1}\} \text{ or } \{\hat{t}_1, \dots, \hat{t}_K\} & \text{if } 0 \leq T - \eta_K < C_d \frac{\sigma^2 \log(\Delta/\alpha)}{\kappa_K^2}, \\ \{\hat{t}_1, \dots, \hat{t}_K\} & \text{otherwise,} \end{cases}$$

where we set $\{\hat{t}_1, \dots, \hat{t}_{K-1}\} = \emptyset$ when $K = 1$. Furthermore,

$$0 \leq \hat{t}_k - \eta_k \leq C_d \frac{\sigma^2 \log(\Delta/\alpha)}{\kappa_k^2}, \quad \forall k = 1, \dots, K,$$

where, if $C = \{\hat{t}_1, \dots, \hat{t}_{K-1}\}$, we define $\hat{t}_K = T$.

The proof of Theorem 3.3 is in Section B and repeatedly uses the proof and results of Theorem 2.2, therefore the proof of Theorem 3.3 holds for any fixed time point T , including random finite stopping times T . In fact, Lemma 3 in Howard et al. (2018) also shown such equivalence in the context of sequential testing.

One of the key steps in the proof is to show that, for any $k \geq 1$,

$$0 \leq \hat{t}_k - \eta_k \leq C_d \frac{\sigma^2 \log(\Delta/\alpha)}{\kappa_k^2} < \Delta/4. \quad (3.2)$$

Once a change point is declared, the procedure is restarted afresh using the latest change point estimate as the new initial time. Due to the last display, we are guaranteed that the new starting point of the procedure, namely $\hat{t}_k$, is only slightly delayed, by an amount no larger than $\Delta/4$ from the ideal starting point η_k . The constant of $1/4$ is not special and could be replaced by any constant smaller than $1/2$. Equivalently, the next change point will be detected starting from $\hat{t}_k$; precisely, we have that

$$\eta_{k+1} - \hat{t}_k = \eta_{k+1} - \eta_k - (\hat{t}_k - \eta_k) > 3\Delta/4.$$

As demonstrated in our proof, this delay is sufficient, under the signal-to-noise condition in [Assumption 3.2](#), to ensure that the procedure is able to detect the next change point η_{k+1} and to localize it with a detection delay no larger than $C_d \sigma^2 \kappa_{k+1}^{-2} \log(\Delta/\alpha)$.

It is important to notice that the quantity K in the theorem statement is itself a random variable. Furthermore, the theorem only guarantees that the procedure will be able to identify, with probability at least $1 - \alpha$ all the K true change point in the interval $\{1, \dots, T\}$ as long as $T - \eta_K \geq C_d \sigma^2 \kappa_K^{-2} \log(\Delta/\alpha)$. If this condition is not satisfied, that is, when the difference between the k th time point and the stopping time is too small in relation to the assumed signal-to-noise ratio [Assumption 3.2](#), there is not guarantee that the procedure will detect the last change point η_K . This is not surprising: The data collected in such a short amount of time between the last-detected change point and T may not in general contain enough information to support that decision.

4. OPTIMALITY

4.1. A Lower Bound on the Detection Delay

In this section, we show that the detection delay bounds derived in [Theorem 2.2](#) are essentially sharp. Toward this end, we adapt to our settings existing techniques for deriving minimax optimality from the literature on change point analysis which, however, only apply asymptotically and under parametric assumption. In particular, Theorem 1 of Lorden (1971) shows that, assuming independent Gaussian instead of sub-Gaussian data,

$$\inf_{\hat{t}} \sup_{\Delta \geq 1} \text{esssup} \mathbb{E}_\Delta \{ (\hat{t} - \Delta)^+ | X_1, \dots, X_\Delta \} \sim \frac{\sigma^2 \log(\gamma)}{\kappa^2}, \quad \text{as } \gamma \rightarrow \infty.$$

where the infimum is taken over all change point estimators $\hat{t}$ such that $\mathbb{E}_\infty(\hat{t}) \geq \gamma$. The last display bears clear similarities with the high-probability delay bounds in [Section 2](#) and, in particular, with the guarantees obtained in [Proposition 2.4](#).

In our next non-asymptotic result, we formally show that the upper bound on the detection delay of [Theorem 2.2](#) is in fact in agreement, with a minimax lower bound on the expected detection delay. The proof adapts arguments used in [Theorem 2](#) of Lai (1998).

Proposition 4.1. Assume that $\{X_i\}_{i=1,2,\dots}$ is a sequence of independent Gaussian random variables satisfying $\mathbb{E}(X_i) = f_i$, $\text{Var}(X_i) = \sigma^2$ and [Assumption 1.2](#). Denote the joint distribution of $\{X_i\}_{i=1,2,\dots}$ as $P_{\kappa,\sigma,\Delta}$. For $\alpha \in (0, 1)$, consider the class of change point estimators

$$\mathcal{D}(\alpha) = \{T : T \text{ is a stopping time with respect to the natural filtration and satisfies } \mathbb{P}_\infty(T < \infty) \leq \alpha\}.$$

Then for all α , κ and σ such that $2\kappa^2\Delta\sigma^{-2} \geq \log(\Delta/\alpha)$ and

$$1 - c_1 > \alpha + \alpha^{1/4}, \quad (3.3)$$

with an absolute constant $c_1 > 0$, it holds that, for any change point time Δ ,

$$\inf_{\hat{t} \in \mathcal{D}(\alpha)} \sup_{P_{\kappa,\sigma,\Delta}} \mathbb{E}_P \left\{ (\hat{t} - \Delta)_+ \right\} \geq \frac{\sigma^2}{2\kappa^2} \log \left(\frac{\Delta}{\alpha} \right). \quad (3.4)$$

Remark 2. The constants in (3.3) have not been optimized and are likely to be improvable.

In the context of multiple change points, the lower bound in [Proposition 4.1](#) further yields that the signal-to-noise ratio condition in [Assumption 3.2](#) is almost necessary to guarantee a consistent change point localization. As remarked after the statement of [Theorem 3.3](#), with probability at least $1 - \alpha$, under this condition, the k th estimated change point is at a distance of no more than $\Delta/4$ from the k th largest true change point, see (3.2). In contrast, whenever [Assumption 3.2](#) is violated in such a way that

$$\kappa^2\Delta/\sigma^2 < \log(\Delta/\alpha)/2, \quad (3.5)$$

then (3.4) implies that

$$\frac{\inf_{\hat{t} \in \mathcal{D}(\alpha)} \sup_{P_{\kappa,\sigma,\Delta}} \mathbb{E}_P \left\{ (\hat{t} - \Delta)_+ \right\}}{\Delta} > 1.$$

Thus the expected detection delay is larger than Δ and, therefore, the true change point cannot be consistently estimated.

Overall, we may conclude that the online multiple change point detection problem is impossible under the low signal-to-noise ratio condition $\kappa^2\Delta/\sigma^2 \log(\Delta/\alpha)$. In contrast, a slightly larger signal-to-noise ratio $\kappa^2\Delta/\sigma^2 \geq \log(\Delta/\alpha)$ is sufficient to guarantee, as shown in [Theorem 3.3](#), that the change points can be accurately estimated with probability no smaller than the prescribed nominal level of $1 - \alpha$. A similar phase transition phenomenon occurs for the offline mean change localization problem, see Wang, Yu, and Rinaldo (2020). In fact, the phase transition boundary is nearly identical in both the online and offline settings.

4.2. Connections with Offline Change Point Detection Problems

A closely related area is the offline change point detection, where one has data $\{X_i\}_{i=1}^T$ and seeks change point estimators $\{\hat{\eta}_k\} \subset \{1, \dots, T\}$. The online and offline change

point analysis shares many similarities. The offline change point results we list below can be found in Wang, Yu, and Rinaldo (2020).

The signal-to-noise ratio and the phase transition. Let Δ be the minimal spacing between two consecutive change points in the offline setting. We remark that in both online and offline problems, the signal-to-noise ratio is of the same form $\kappa^2 \Delta / \sigma^2$. In both problems, the parameter spaces can be partitioned into feasibility and infeasibility regimes by this signal-to-noise ratio. In Wang, Yu, and Rinaldo (2020), it is shown that in the univariate offline change point detection problem, both the lower and upper bounds on the signal-to-noise ratio are of order $\log(T)$.

When deriving the estimation error, since one has collected all the data in advance in the offline setting, the signal-to-noise ratio is lower bounded by $\log(T)$ and as a result, the estimation error only depends on the model parameters. This is not the case in the online setting, where the total number of data points examined is also a random variable. In this case, additional information is needed. In Theorem 2.2, we choose to control the upper bound of the type I error α . As a result, the estimation error, that is, the detection delay, is a function also of α .

5. EXPERIMENTS

To validate the theoretical results presented in this paper, we consider the empirical performance of Algorithm 1 with choices of $\{b_t\}$ guided by Theorem 2.2. Since the main focus of our theoretical results lies in the theoretical optimality in terms of rates, we sacrifice the sharpness in terms of constants for presentational convenience. It can be seen from Appendix A that, we have used more conservative $\{b_t\}$ in Theorem 2.2. To demonstrate the numerical performances, guided by the results in Appendix A, we consider the thresholds

$$\sigma \sqrt{4 \log \left\{ \frac{2t^2}{s(t-s)} \right\} - 2 \log(\alpha)}$$

in Algorithm 1.

As competitors we consider the thresholding rule from Theorem 6 in Maillard (2019) (M). We also compare against Chen (2019). For comparing with the latter, we consider its three different statistics, including the “original” (ORI) which specifies the original edge-count scan statistic, the weighted edge-count scan statistic (W), and the generalized edge-count scan statistic (G). These statistics are calibrated using 300 Monte Carlo simulations where in each simulation we generate a sequence $\tilde{y}_1, \dots, \tilde{y}_{100} \sim \mathcal{N}(0, \sigma^2)$. All our comparisons with Chen (2019) are done with the R package gStream (Chen and Chu 2019).

With regard to generating the data, we set $T = 400$ and construct $y \in \mathbb{R}^T$ as

$$y_t \sim \begin{cases} \mathcal{N}(0, \sigma^2), & t < \Delta, \\ \mathcal{N}(1, \sigma^2), & t \geq \Delta, \end{cases}$$

with $\Delta \in \{50, \infty\}$ and $\sigma \in \{0.1, 0.5, 1, 0.8, 1.2\}$.

Table 1. Numerical results.

α	σ	Method	PFA	DD	α	σ	Method	PFA	DD
0.05	0.1	Alg 1	0.04	0.00	0.1	0.1	Alg 1	0.00	0.00
0.05	0.1	M	0.02	0.00	0.1	0.1	M	0.00	0.00
0.05	0.1	ORI	0.06	7.89	0.1	0.1	ORI	0.10	7.28
0.05	0.1	W	0.08	7.46	0.1	0.1	W	0.10	6.82
0.05	0.1	G	0.04	7.73	0.1	0.1	G	0.06	7.11
0.05	0.5	Alg 1	0.04	4.81	0.1	0.5	Alg 1	0.07	4.60
0.05	0.5	M	0.00	5.33	0.1	0.05	M	0.02	4.92
0.05	0.5	ORI	0.04	17.10	0.1	0.05	ORI	0.12	15.65
0.05	0.5	W	0.02	17.32	0.1	0.05	W	0.08	16.67
0.05	0.5	G	0.04	18.11	0.1	0.05	G	0.08	16.78
0.05	0.8	Alg 1	0.06	11.77	0.1	0.8	Alg 1	0.12	9.53
0.05	0.8	M	0.02	15.59	0.1	0.8	M	0.10	14.24
0.05	0.8	ORI	0.05	120.48	0.1	0.8	ORI	0.10	94.25
0.05	0.8	W	0.03	145.13	0.1	0.8	W	0.07	87.28
0.05	0.8	G	0.03	152.29	0.1	0.8	G	0.10	99.37
0.05	1.2	Alg 1	0.05	27.16	0.1	1.2	Alg 1	0.06	23.08
0.05	1.2	M	0.01	80.09	0.1	1.2	M	0.03	66.52
0.05	1.2	ORI	0.03	241.10	0.1	1.2	ORI	0.08	209.31
0.05	1.2	W	0.02	249.46	0.1	1.2	W	0.13	237.14
0.05	1.2	G	0.03	253.77	0.1	1.2	G	0.12	226.48

For both metrics the numbers are obtained by avergaing over 100 Monte Carlo simulations.
PFA, probability of false alarm; DD, detection delay.

As for the performance of different methods, we consider the choices $\alpha \in \{0.05, 0.1\}$. Then for each choice of α and with $\nu = \infty$, we report the proportion of false alarms (PFA) in 100 Monte Carlo trials, where we run each competing method up to $T = 400$. Furthermore, when $\Delta = 50$, denoting by $\hat{t}_i$ the output of a method across the $i = 1, 2, \dots, 100$ Monte Carlo trials, we report the detection delay defined as

$$\text{DD} = \frac{1}{100} \sum_{i=1}^{1000} (\min\{\hat{t}_i, T\} - \Delta)_+. \quad (5.1)$$

The results in Table 1 show that, overall, the best performance is attained by Algorithm 1 followed by Maillard (2019). This is reflected in both PFA and DD.

6. DISCUSSION

We have investigated various aspects of online mean change point estimation with sub-Gaussian, independent variates. In our analysis, we have adopted finite sample settings and allowed for the model parameters to vary so as to express all the sources of statistical hardness in the problem. We have presented simple but effective algorithms for online change point localization based on classical CUSUM process that are provably nearly optimal. We have also identified combinations of model parameters for which the change point estimation task is impossible and showed that our procedure succeeds over nearly all the other combinations, including those for which the problem is most difficulty and barely solvable. Interestingly, the online estimation rates and phase transition boundaries we determine are nearly identical to those holding in the offline settings; see Wang, Yu, and Rinaldo (2020). The framework we established in this paper

can be used to study different high-dimensional problems. This will be left as future work.

In order to control the fluctuations of the CUSUM statistics, we have adopted two different approaches. [Theorems 2.2](#) and [3.3](#) rely on peeling arguments, as detailed in [Lemmas A.1](#) and [A.2](#). In contrast, the results in [Section 2.2](#) rely on simpler union bound arguments. Interestingly, even using peeling, we were only able to establish that the CUSUM statistic process

$$t \mapsto \sup_{t \geq 2} \max_{1 \leq s < t} \left| \left\{ \frac{s}{t(t-s)} \right\}^{1/2} \sum_{l=s+1}^t (X_l - f_l) \right|,$$

has fluctuations of order $\sqrt{\{\log(t) + \log(\alpha)\}/t}$. One may have expected to see a better dependence on t , such as in $\sqrt{\log \log(t)}$. However, the above CUSUM process does not have a clear martingale structure. As a result, peeling techniques only yield crossing boundaries that scale like $\sqrt{\log(t)}$, albeit with better constants than those of the same order stemming from union bound arguments. Since we have not made attempts to optimize constants, both approaches—peeling and union bounds (or Bonferroni-type arguments)—deliver the same localization rates in our settings.

DISCLOSURE

The authors have no conflicts of interest to report.

APPENDICES

We include all the technical details and some additional simulation results in the Appendices.

APPENDIX A. CONCENTRATION INEQUALITIES

Lemma A.1. *For any $\alpha > 0$, it holds that*

$$\mathbb{P} \left\{ \exists s, t \in \mathbb{N}, \ t > 1, \ s \in [1, t) : \left| \left(\frac{t-s}{ts} \right)^{1/2} \sum_{l=1}^s (X_l - f_l) - \left\{ \frac{s}{t(t-s)} \right\}^{1/2} \sum_{l=s+1}^t (X_l - f_l) \right| > 2^{3/2} \sigma \log^{1/2}(t/\alpha) \right\} \leq \alpha.$$

Proof. It holds that for any sequence $\{\varepsilon_t > 0\}$,

$$\mathbb{P} \left\{ \exists s, t \in \mathbb{N}, \ t > 1, \ s \in [1, t) : \left| \left(\frac{t-s}{ts} \right)^{1/2} \sum_{l=1}^s (X_l - f_l) - \left\{ \frac{s}{t(t-s)} \right\}^{1/2} \sum_{l=s+1}^t (X_l - f_l) \right| \geq \varepsilon_t \right\}$$

$$\begin{aligned}
&\leq \sum_{j=1}^{\infty} \mathbb{P} \left\{ \max_{2^j \leq t < 2^{j+1}} \max_{1 \leq s < t} \left| \left(\frac{t-s}{ts} \right)^{1/2} \sum_{l=1}^s (X_l - f_l) - \left\{ \frac{s}{t(t-s)} \right\}^{1/2} \sum_{l=s+1}^t (X_l - f_l) \right| \geq \varepsilon_t \right\} \\
&\leq \sum_{j=1}^{\infty} 2^j \max_{2^j \leq t < 2^{j+1}} \mathbb{P} \left\{ \max_{1 \leq s < t} \left| \left(\frac{t-s}{ts} \right)^{1/2} \sum_{l=1}^s (X_l - f_l) - \left\{ \frac{s}{t(t-s)} \right\}^{1/2} \sum_{l=s+1}^t (X_l - f_l) \right| \geq \varepsilon_t \right\} \\
&\leq \sum_{j=1}^{\infty} 2^j \max_{2^j \leq t < 2^{j+1}} t \mathbb{P} \{ |W| \geq \varepsilon_t \} \leq \sum_{j=1}^{\infty} 2^{2j+1} \max_{2^j \leq t < 2^{j+1}} \mathbb{P} \{ |W| \geq \varepsilon_t \},
\end{aligned}$$

where W is a mean zero sub-Gaussian random variable with $\|W\|_{\psi_2} \leq \sigma$. Now put

$$\varepsilon_t = \sqrt{2}\sigma [2 \log(t) + 3 \log \log(t) + \log \{ \log(t) + \log(2) \} - 2 \log \log(2) - \log(\alpha)]^{1/2}.$$

Due to the sub-Gaussianity, we have that for any $\zeta > 0$, $\mathbb{P}\{|W| \geq \zeta\} < 2 \exp(-2^{-1}\zeta^2/\sigma^2)$ (e.g., (2.9) in Wainwright 2019). Therefore,

$$\begin{aligned}
&\mathbb{P} \left\{ \exists s, t \in \mathbb{N}, \quad t > 1, \quad s \in [1, t) : \left| \left(\frac{t-s}{ts} \right)^{1/2} \sum_{l=1}^s (X_l - f_l) \right. \right. \\
&\quad \left. \left. - \left\{ \frac{s}{t(t-s)} \right\}^{1/2} \sum_{l=s+1}^t (X_l - f_l) \right| \geq \varepsilon_t \right\} \\
&\leq \sum_{j=1}^{\infty} \max_{2^j \leq t < 2^{j+1}} \exp[(2j+2) \log(2) - 2 \log(t) - 3 \log \log(t) - \log \{ \log(t) + \log(2) \} \\
&\quad + 2 \log \log(2) + \log(\alpha)] \\
&\leq \sum_{j=1}^{\infty} \exp \{ (2j+2) \log(2) - 2j \log(2) - 3 \log(j) - \log \{ (j+1) \log(2) \} \\
&\quad + 2 \log \log(2) + \log(\alpha) \} \\
&\leq \alpha \sum_{j=1}^{\infty} \frac{2^{2j+2} \log^2 2}{2^{2j} j^3 (j+1) \log 2} \leq \alpha \sum_{j=1}^{\infty} \frac{1}{j(j+1)} \leq \alpha.
\end{aligned}$$

For simplicity, we let

$$\varepsilon_t = 2^{3/2} \sigma \log^{1/2}(t/\alpha),$$

which satisfies that for any $t \geq 2$ and $\alpha \in (0, 1)$,

$$2^{3/2} \log^{1/2}(t/\alpha) \geq \sqrt{2} [2 \log(t) + \log \log(t) + \log \{ \log(t) + \log(2) \} - 2 \log \log(2) - \log(\alpha)]^{1/2}. \quad (\text{A.1})$$

This completes the proof. $\square$

Lemma A.2. For any $\alpha > 0$, it holds that

$$\begin{aligned}
&\mathbb{P} \left\{ \exists e, s, t \in \mathbb{N}, \quad e \geq 0, \quad s > e, \quad t > s : \left| \left\{ \frac{t-s}{(s-e)(t-e)} \right\}^{1/2} \sum_{l=e+1}^s (X_l - f_l) \right. \right. \\
&\quad \left. \left. - \left\{ \frac{s-e}{(t-s)(t-e)} \right\}^{1/2} \sum_{l=s+1}^t (X_l - f_l) \right| > 4\sigma \log^{1/2}(t/\alpha) \right\} \leq \alpha.
\end{aligned}$$

Proof. For any integer triplet (e, s, t) , $0 \leq e < s < t$, any sequence $\{\varepsilon_{e,s,t} > 0\}$, let

$$\begin{aligned} \mathcal{G}(e, s, t) &= \left\{ \left| \left\{ \frac{t-s}{(s-e)(t-e)} \right\}^{1/2} \sum_{l=e+1}^s (X_l - f_l) \right. \right. \\ &\quad \left. \left. - \left\{ \frac{s-e}{(t-s)(t-e)} \right\}^{1/2} \sum_{l=s+1}^t (X_l - f_l) \right| > \varepsilon_{e,s,t} \right\} \\ &= \{G(e, s, t) > \varepsilon_{e,s,t}\}. \end{aligned}$$

We have that

$$\begin{aligned} &\mathbb{P}\{\exists e, s, t \in \mathbb{N}, e \geq 0, s > e, t > s : \mathcal{G}(e, s, t)\} \\ &\leq \mathbb{P}\{\exists s, t \in \mathbb{N}, e = 0, s > e, t > s : \mathcal{G}(0, s, t)\} \\ &\quad + \mathbb{P}\{\exists s, t \in \mathbb{N}, e \geq 1, s > e, t > s : \mathcal{G}(e, s, t)\} \\ &= (I) + (II). \end{aligned}$$

As for (I), it follows from Lemma A.1 that, with $\varepsilon_{0,s,t}$ defined as

$$\begin{aligned} \varepsilon_{0,s,t} = \varepsilon_t^{(1)} &= \sqrt{2}\sigma[2\log(t) + 3\log\log(t) + \log\{\log(t) + \log(2)\} \\ &\quad - 2\log\log(2) + \log(2) - \log(\alpha)]^{1/2}, \end{aligned}$$

we have that

$$(I) \leq \alpha/2.$$

As for (II), we have that

$$\begin{aligned} (II) &\leq \sum_{j=1}^{\infty} \mathbb{P}\left\{ \max_{2^j \leq t < 2^{j+1}} \max_{1 \leq e < t-1} \max_{e < s < t} G(e, s, t) > \varepsilon_{e,s,t} \right\} \\ &\leq \sum_{j=1}^{\infty} \sum_{m=0}^{j-1} \mathbb{P}\left\{ \max_{2^j \leq t < 2^{j+1}} \max_{2^m \leq e < 2^{m+1}} \max_{e < s < t} G(e, s, t) > \varepsilon_{e,s,t} \right\} \\ &\quad + \sum_{j=1}^{\infty} \mathbb{P}\left\{ \max_{2^j \leq t < 2^{j+1}} \max_{2^j \leq e < t} \max_{e < s < t} G(e, s, t) > \varepsilon_{e,s,t} \right\} \\ &< \sum_{j=1}^{\infty} \left(\sum_{m=0}^{j-1} 2^j 2^m 2^{j+1} + 2^j 2^j 2^j \right) \max_{\substack{e < s < t \\ 2^j \leq t < 2^{j+1}}} \mathbb{P}\{\mathcal{G}(e, s, t)\} \\ &\leq \sum_{j=1}^{\infty} 2^{3j+2} \max_{\substack{e < s < t \\ 2^j \leq t < 2^{j+1}}} \mathbb{P}\{\mathcal{G}(e, s, t)\}. \end{aligned}$$

For any $e > 0$, let

$$\begin{aligned} \varepsilon_{e,s,t} = \varepsilon_t^{(2)} &= \sqrt{2}\sigma[3\log(t) + 3\log\log(t) + \log\{\log(t) + \log(2)\} \\ &\quad + 3\log(2) - 2\log\log(2) - \log(\alpha)]^{1/2}. \end{aligned}$$

Due to the sub-Gaussianity, we have that

$$(II) \leq \alpha/2.$$

For simplicity, we let

$$\varepsilon_t = 4\sigma \log^{1/2}(t/\alpha),$$

such that

$$\begin{aligned} 4\log^{1/2}(t/\alpha) &\geq 2^{1/2}[3\log(t) + 3\log\log(t) + \log\{\log(t) + \log(2)\} \\ &\quad + 3\log(2) - 2\log\log(2) - \log(\alpha)]^{1/2}. \end{aligned} \tag{A.2}$$

We therefore completes the proof. $\square$

APPENDIX B. PROOFS OF THE MAIN RESULTS

Proof of Theorem 2.2. This proof consists of four steps. In Step 1 we define a large probability event, where the rest of the proof is conducted. In Step 2 we conclude the claims on type I error controls. In Step 3 we change the detection delay upper bounding problem to upper bounding the delay with respect to an easier to handle quantity t^* . In Step 4 we conclude the proof.

Step 1. Define the event

$$\mathcal{B} = \left\{ \forall s, t \in \mathbb{N}, t > 1, s \in [1, t) : \left| \left(\frac{t-s}{ts} \right)^{1/2} \sum_{l=1}^s (X_l - f_l) - \left\{ \frac{s}{t(t-s)} \right\}^{1/2} \sum_{l=s+1}^t (X_l - f_l) \right| < b_t \right\}.$$

It follows from Lemma A.1 that $\mathbb{P}\{\mathcal{B}\} > 1 - \alpha$. Throughout the proof we assume that the event $\mathcal{B}$ holds.

For any $s, t \in \mathbb{N}, 1 \leq s < t$, it holds that $|\widehat{D}_{s,t} - D_{s,t}| < b_t$, which implies that

$$D_{s,t} + b_t > \widehat{D}_{s,t} > D_{s,t} - b_t. \quad (\text{B.1})$$

Step 2. For any $t \leq \Delta$, we have that $D_{s,t} = 0$, for all $s \in [1, t)$. Thus, using (B.1), we conclude that, $\widehat{t} > t$ and, therefore that $\widehat{t} > \Delta$. We have now shown (2.2) and (2.3).

Step 3. Now we consider any $t > \Delta$. If there exists $s \in [1, t)$ such that $\widehat{D}_{s,t} > b_t$, then $d = \widehat{t} - \Delta \leq t - \Delta$. Thus, $d \leq \tilde{t} - \Delta$, where

$$\tilde{t} = \min\{t > \Delta, \exists s \in [\Delta, t), \widehat{D}_{s,t} > b_t\},$$

and any upper bound on $\tilde{t}$ will also be an upper bound on d . Thus, our task becomes that of computing a sharp upper bound on $\tilde{t}$. To that effect, notice that, when $\Delta \leq s < t$,

$$D_{s,t} = \Delta \left(\frac{t-s}{ts} \right)^{1/2} |\mu_1 - \mu_2| = \Delta \left(\frac{t-s}{ts} \right)^{1/2} \kappa,$$

and, because of (B.1) again,

$$\widehat{D}_{s,t} \geq \Delta \left(\frac{t-s}{ts} \right)^{1/2} \kappa - b_t.$$

As a result, we obtain that $\tilde{t} \leq t^*$, where

$$t^* = \min \left\{ t > \Delta : \max_{s \in [\Delta, t)} \left\{ \Delta \left(\frac{t-s}{ts} \right)^{1/2} \kappa - 2b_t \right\} \geq 0 \right\}.$$

Step 4. Write for convenience $m = t^* - \Delta$, so that $d \leq m$. Recalling that $b_t = 2^{3/2} \sigma \log^{1/2}(t/\alpha)$, we seek the smallest integer m such that

$$\max_{s \in [\Delta, m+\Delta)} \left[\Delta \kappa \left\{ \frac{m+\Delta-s}{(m+\Delta)s} \right\}^{1/2} - 2^{5/2} \sigma \log^{1/2}\{(m+\Delta)/\alpha\} \right] > 0,$$

which is equivalent to finding the smallest integer m such that

$$\max_{s \in [\Delta, m+\Delta)} \left[\Delta^2 \kappa^2 - 32 \sigma^2 \frac{s(m+\Delta)}{m+\Delta-s} \log \{(m+\Delta)/\alpha\} \right] > 0.$$

In turn, the above task corresponds to that of computing the smallest integer m such that

$$\begin{aligned} \Delta^2 \frac{\kappa^2}{\sigma^2} &> \min_{s \in [\Delta, m+\Delta)} \left[32 \frac{s(m+\Delta)}{m+\Delta-s} \log \{(m+\Delta)/\alpha\} \right] \\ &= 32 \frac{\Delta(m+\Delta)}{m} \log \{(m+\Delta)/\alpha\}, \end{aligned}$$

or, equivalently, such that

$$m \left\lceil \frac{\Delta \kappa^2}{32 \sigma^2} - \log \{(m + \Delta)/\alpha\} \right\rceil > \Delta \log \{(m + \Delta)/\alpha\}, \quad (\text{B.2})$$

under [Assumption 2.1](#).

Let C_d be an absolute constant large enough and also upper bounded by C_{SNR} . The claimed result now follows once we show that the value

$$m^* = \lceil C_d \log(\Delta/\alpha) \sigma^2 \kappa^{-2} \rceil$$

satisfies (B.2). To see this, assume for simplicity that $C_d \log(\Delta/\alpha) \sigma^2 \kappa^{-2}$ is an integer; if not, the proof only requires trivial modifications. We first point out that $m^* \leq \Delta$ because of [Assumption 2.1](#) and the fact that $C_d \leq C_{\text{SNR}}$. Now, the lefthand side of inequality (B.2) is equal, for this choice of m , to

$$C_d \log(\Delta/\alpha) \frac{\Delta}{32} - C_d \frac{\sigma^2}{\kappa^2} \log(\Delta/\alpha) \log \left\{ \frac{C_d \log(\Delta/\alpha) \sigma^2 / \kappa^2 + \Delta}{\alpha} \right\}. \quad (\text{B.3})$$

Using again [Assumption 2.1](#) and the fact that $m^* \leq \Delta$, the second term in the previous expression is upper bounded by

$$\frac{2C_d}{C_{\text{SNR}}} \Delta \log(\Delta/\alpha),$$

due to the fact that $2 \log(x) \geq \log(2x)$, $x \geq 2$. Thus, the quantity in (B.3) is lower bounded by

$$\Delta \log(\Delta/\alpha) \left(\frac{C_d}{32} - \frac{2C_d}{C_{\text{SNR}}} \right) \geq 2\Delta \log(\Delta/\alpha) \geq \Delta \log(2\Delta/\alpha) \geq \Delta \log((m^* + \Delta)/\alpha),$$

where the first inequality is justified by first choosing a large enough C_d and then choosing C_{SNR} larger than C_d , and the second and third inequalities follow from $\log(\Delta/\alpha) \geq 0$ and $m^* \leq \Delta$, respectively. Thus, combining the last display with (B.2) and (B.3) yields (2.4). $\square$

Proof of Proposition 2.3. It follows from Step 1 in the proof of [Theorem 2.2](#) that on the event $\mathcal{B}$, it holds that

$$\widehat{D}_{s,t} < D_{s,t} + 2^{3/2} \sigma \log^{1/2}(t/\alpha), \quad 1 \leq s < t.$$

It follows from Step 2 in the proof of [Theorem 2.2](#) that we only need to consider $t > \Delta$. This leaves us two situations: $s \leq \Delta$ and $s \geq \Delta$. In fact in both these two situations, one only needs to deal with the case $s = \Delta$, therefore we only show $s \geq \Delta$ here.

When $\Delta \leq s < t$, we have $D_{s,t} = \kappa \Delta \{(t-s)/ts\}^{1/2}$ and therefore on the event $\mathcal{B}$, we have that

$$\begin{aligned} \max_{\Delta \leq s < t} \widehat{D}_{s,t} &\leq \max_{\Delta \leq s < t} \kappa \Delta \left(\frac{t-s}{ts} \right)^{1/2} + 2^{3/2} \sigma \log^{1/2}(t/\alpha) \\ &\leq \kappa \Delta \left(\frac{t-s}{ts} \right)^{1/2} + 2^{3/2} \sigma \log^{1/2}(t/\alpha) \leq \kappa \Delta^{1/2} + 2^{3/2} \sigma \log^{1/2}(t/\alpha) < b_t, \end{aligned}$$

where the last inequality follows from (2.6). We therefore completes the proof. $\square$

Proof of Proposition 2.4. This proof consists of three steps. In **Step 1** we define a large probability event, where the rest of the proof is conducted. We also conclude the control on the average run length. In **Step 2** we change the detection delay upper bounding problem to upper bounding the delay with respect to an easier to handle quantity t^* . In **Step 3** we conclude the proof.

Step 1. Define the event

$$\mathcal{C} = \left\{ t \in \{2, \dots, \gamma + 2\}, s \in \mathbb{N} \cap [1, t) : \left| \left(\frac{t-s}{ts} \right)^{1/2} \sum_{l=1}^s (X_l - f_l) - \left\{ \frac{s}{t(t-s)} \right\}^{1/2} \sum_{l=s+1}^t (X_l - f_l) \right| < b_t \right\}.$$

We have that

$$\begin{aligned} 1 - \mathbb{P}\{\mathcal{C}\} &\leq (\gamma + 1) \max_{t \in \{2, \dots, \gamma+2\}} \mathbb{P}\left\{ \exists s \in \mathbb{N} \cap [1, t) : \left| \left(\frac{t-s}{ts} \right)^{1/2} \sum_{l=1}^s (X_l - f_l) - \left\{ \frac{s}{t(t-s)} \right\}^{1/2} \sum_{l=s+1}^t (X_l - f_l) \right| \geq b_t \right\} \\ &\leq (\gamma + 1)^2 \mathbb{P}\{|W| > b_t\} \leq 2(\gamma + 1)^2 \exp\left(-\frac{b_t^2}{2\sigma^2}\right) = (\gamma + 1)^{-1}, \end{aligned}$$

where W is a mean zero sub-Gaussian random variable with $\|W\|_{\psi_2} \leq \sigma$ and the Hoeffding inequality (e.g., (2.9) in Wainwright 2019).

Therefore we have that

$$\mathbb{E}_\infty(\hat{t}) = \sum_{t=2}^{\infty} \mathbb{P}\{\hat{t} \geq t\} \geq \sum_{t=2}^{\gamma+2} \mathbb{P}\{\hat{t} \geq t\} \geq (\gamma + 1) \mathbb{P}\{\hat{t} > \gamma + 2\} \geq (\gamma + 1) \left(1 - \frac{1}{\gamma + 1}\right) = \gamma.$$

Step 2. We have that

$$\mathbb{P}\{\hat{t} > \gamma + 2\} = \mathbb{P}\{\mathcal{C}\} \geq 1 - (\gamma + 1)^{-1},$$

When $\gamma > \max\{\Delta - 2, 1\}$, it holds that

$$\mathbb{P}\{d > 0\} \geq 1 - (\gamma + 1)^{-1}.$$

For any $t > \Delta$, if there exists $s \in [1, t)$ such that $\widehat{D}_{s,t} > b_t$, then $d \leq t - \Delta$. It suffices to find

$$\tilde{t} = \min\{t : t > \Delta, \exists s \in (\Delta, t), \widehat{D}_{s,t} > b_t\},$$

and any upper bound on $\tilde{t}$ will also be an upper bound on d , when the signal-to-noise ratio condition in (2.8) is satisfied. Then, our task becomes that of computing a sharp upper bound on $\tilde{t}$. To that effect, notice that, when $\Delta \leq s < t$,

$$D_{s,t} = \Delta \left(\frac{t-s}{ts} \right)^{1/2} |\mu_1 - \mu_2| = \Delta \left(\frac{t-s}{ts} \right)^{1/2} \kappa,$$

and on the event $\mathcal{C}$,

$$\widehat{D}_{s,t} \geq \Delta \left(\frac{t-s}{ts} \right)^{1/2} \kappa - b_t.$$

As a result, we obtain that $\tilde{t} \leq t^*$, where

$$t_* = \min \left\{ t > \Delta : \max_{s \in (\Delta, t)} \left\{ \Delta \left(\frac{t-s}{ts} \right)^{1/2} \kappa - 2b_t \right\} \geq 0 \right\}.$$

Step 3. Write for convenience $m = t - \Delta$, so that $d \leq m$. Recalling that

$$b_t = 6^{1/2} \sigma \log^{1/2} \left\{ 2^{1/3} (\gamma + 1) \right\},$$

we seek the smallest integer m such that

$$\max_{s \in [\Delta, m+\Delta]} \left[\Delta \kappa \left\{ \frac{m+\Delta-s}{(m+\Delta)s} \right\}^{1/2} - 2 \times 6^{1/2} \sigma \log^{1/2} \{2^{1/3}(\gamma+1)\} \right] > 0,$$

which is equivalent to finding the smallest integer m such that

$$\max_{s \in [\Delta, m+\Delta]} \left\{ \Delta^2 \kappa^2 - 24\sigma^2 \frac{s(m+\Delta)}{m+\Delta-s} \log \{2^{1/3}(\gamma+1)\} \right\} > 0.$$

In turn, the above task corresponds to that of computing the smallest integer m such that

$$\begin{aligned} t\Delta^2 \kappa^2 &> \min_{s \in [\Delta, m+\Delta]} \left\{ 24\sigma^2 \frac{s(m+\Delta)}{m+\Delta-s} \log \{2^{1/3}(\gamma+1)\} \right\} \\ &= 24\sigma^2 \frac{\Delta(m+\Delta)}{m} \log \{2^{1/3}(\gamma+1)\}, \end{aligned}$$

or, equivalently, such that

$$m \left[\frac{\Delta \kappa^2}{24\sigma^2} - \log \{2^{1/3}(\gamma+1)\} \right] > \Delta \log \{2^{1/3}(\gamma+1)\}. \quad (\text{B.4})$$

Let C'_d be a large enough absolute constant. The claimed result now follows once we show that the value

$$m^* = \lceil C'_d \sigma^2 \kappa^{-2} \log \{2^{1/3}(\gamma+1)\} \rceil$$

satisfies (B.4). To see this, assume for simplicity that $C'_d \sigma^2 \kappa^{-2} \log \{2^{1/3}(\gamma+1)\}$ is an integer; if not, the proof only requires trivial modifications. Now, the lefthand side of inequality (B.4) is equal, for this choice of m , to

$$C'_d \log \{2^{1/3}(\gamma+1)\} \frac{\Delta}{24} - C'_d \sigma^2 \kappa^{-2} \log^2 \{2^{1/3}(\gamma+1)\}. \quad (\text{B.5})$$

Using again (2.8), the second term in the previous expression is upper bounded by

$$\frac{2C'_d}{C_{\text{SNR}}} \Delta \log \{2^{1/3}(\gamma+1)\},$$

due to the fact that $2 \log(x) \geq \log \{2^{1/3}(x+1)\}$, $x \geq 2$. Thus, the quantity in (B.5) is lower bounded by

$$\Delta \log \{2^{1/3}(\gamma+1)\} \left(\frac{C'_d}{24} - \frac{2C'_d}{C_{\text{SNR}}} \right) \geq 2\Delta \log \{2^{1/3}(\gamma+1)\} \geq \Delta \log \{2^{1/3}(\gamma+1)\},$$

where the first inequality is justified with first choosing a large enough C_d , then a large enough C_{SNR} . Thus we conclude the proof with a large enough absolute constant $C_d > 0$. $\square$

Proof of Proposition 2.5. When $\Delta = \infty$, assuming that $f_1 = f_2 = \dots$ and letting $e_{\gamma,t} = \max\{t - \gamma + 1, 1\}$, we have that

$$\begin{aligned} &\sup_{v>1} \mathbb{P}(v \leq \hat{t} \leq v + \gamma) \\ &\leq \sup_{v>1} \mathbb{P} \left\{ \exists s \in \mathbb{N} \cap (v, v + \gamma), t \in N \cap [e_{\gamma,t}, t-1] : \hat{D}_{e_{\gamma,t}, s, t} > b_\gamma \right\} \\ &\leq \gamma^2 \mathbb{P}\{|W| > b_\gamma\} \leq \alpha, \end{aligned}$$

where W is a mean zero sub-Gaussian random variable with $\|W\|_{\psi_2} \leq \sigma$ and satisfies that any $\zeta > 0$, $\mathbb{P}\{|W| \geq \zeta\} < 2 \exp(-2^{-1} \zeta^2 / \sigma^2)$.

When Assumption 1.2 is imposed, we consider $t > \Delta$. For any $\Delta < t < \Delta + \gamma/2$, if there exists $s \in [1, t)$ such that $\hat{D}_{e_{\gamma,s}, t} > b_\gamma$, then $d \leq t - \Delta$. It suffices to find

$$\tilde{t} = \min\{t : \Delta < t < \Delta + \gamma/2, \hat{D}_{t-\gamma, \Delta, t} > b_\gamma\},$$

and any upper bound on $\tilde{t}$ will also be an upper bound on d , when (2.9) is satisfied. Thus, our task becomes that of computing a sharp upper bound on $\tilde{t}$. To that effect, notice that, when $\Delta = s < t < \Delta + \gamma/2$, we have

$$D_{t-\gamma, \Delta, t} = \left\{ \frac{(t-\Delta)(\Delta-t+\gamma)}{\gamma} \right\}^{1/2} |\mu_1 - \mu_2| = \left\{ \frac{(t-\Delta)(\Delta-t+\gamma)}{\gamma} \right\}^{1/2} \kappa.$$

In addition, it holds that

$$\mathbb{P}\{\mathcal{G}\} = \mathbb{P}\left\{\exists t \in (\Delta, \Delta + \gamma/2) : |\widehat{D}_{t-\gamma, \Delta, t} - D_{t-\gamma, \Delta, t}| > b_\gamma\right\} \leq \gamma/2\alpha\gamma^{-2} < \alpha.$$

Therefore on the event $\mathcal{G}$,

$$\widehat{D}_{t-\gamma, \Delta, t} \geq \left\{ \frac{(t-\Delta)(\Delta-t+\gamma)}{\gamma} \right\}^{1/2} \kappa - b_\gamma.$$

As a result, we obtain that $\tilde{t} \leq t^*$, where

$$t_* = \min \left\{ t \in (\Delta, \Delta + \gamma/2) : \left\{ \frac{(t-\Delta)(\Delta-t+\gamma)}{\gamma} \right\}^{1/2} \kappa - 2b_\gamma \geq 0 \right\}.$$

Write for convenience $m = t_* - \Delta$, so that $d \leq m$. Due to the form of b_γ , we seek the smallest integer m such that

$$\left\{ \frac{m(\gamma-m)}{\gamma} \right\}^{1/2} \kappa - 2^{3/2} \sigma \log^{1/2}(2\gamma^2 \alpha^{-1}) > 0.$$

We now to show that $m = \lceil C_d \sigma^2 \log(\gamma/\alpha) \kappa^{-2} \rceil$ satisfies the above. Due to (2.9), we have that $(\gamma-m)/\gamma \geq 1/2$, then

$$\frac{m(\gamma-m)}{\gamma} \kappa^2 \geq \kappa^2 m/2 > 2^{3/2} \sigma \log^{1/2}(2\gamma^2 \alpha^{-1}),$$

which completes the proof. □

Proof of Corollary 2.6. This proof consists of four steps. In **Step 1** we define a large probability event, where the rest of the proof is conducted. In **Step 2** we conclude the claims on type I error controls. In **Step 3** we change the detection delay upper bounding problem to upper bounding the delay with respect to an easier to handle quantity t^* . In Step 4 we conclude the proof.

Step 1. Define the event

$$\mathcal{B} = \left\{ \forall s, t \in \mathbb{N}, t > 1, s \in \{t-2^0, \dots, t-2^{\lfloor \log(t)/\log(2) \rfloor}\} : \left| \left(\frac{t-s}{ts} \right)^{1/2} \sum_{l=1}^s (X_l - f_l) - \left\{ \frac{s}{t(t-s)} \right\}^{1/2} \sum_{l=s+1}^t (X_l - f_l) \right| < b_t \right\}.$$

It follows from Lemma A.1 that $\mathbb{P}\{\mathcal{B}\} > 1 - \alpha$. Throughout the proof we assume that the event $\mathcal{B}$ holds.

For any $s, t \in \mathbb{N}, 1 \leq s < t$, it holds that $|\widehat{D}_{s,t} - D_{s,t}| < b_t$, which implies that

$$D_{s,t} + b_t > \widehat{D}_{s,t} > D_{s,t} - b_t. \quad (\text{B.6})$$

For any $t \geq 1$, we have that

$$\mathcal{S}(t) = \{t-2^0, \dots, t-2^{\lfloor \log(t)/\log(2) \rfloor}\}.$$

Step 2. For any $t \leq \Delta$, we have that $D_{s,t} = 0$, for all $s \in \mathcal{S}(t)$. Using (25), we conclude that $\widehat{t} > \Delta$ and therefore that $\widehat{t} > \Delta$.

Step 3. Now we consider $t > \Delta$. If there exists $s \in \mathcal{S}(t)$ such that $\widehat{D}_{s,t} > b_t$, then $d \leq t - \Delta$. Thus, $d \leq \tilde{t} - \Delta$, where

$$\tilde{t} = \min\{t > \Delta, \exists s \in [\Delta, t), \hat{D}_{s,t} > b_t\},$$

and any upper bound on $\tilde{t}$ will also be an upper bound on d , when the signal-to-noise ratio constraint specified by [Assumption 2.1](#) is satisfied. Thus, our task becomes that of computing a sharp upper bound on $\tilde{t}$. To that effect, notice that, when $\Delta \leq s < t$,

$$D_{s,t} = \Delta \left(\frac{t-s}{ts} \right)^{1/2} |\mu_1 - \mu_2| = \Delta \left(\frac{t-s}{ts} \right)^{1/2} \kappa,$$

and, because of [\(B.6\)](#) again,

$$\hat{D}_{s,t} \geq \Delta \left(\frac{t-s}{ts} \right)^{1/2} \kappa - b_t.$$

As a result, we obtain that $\tilde{t} \leq t^*$, where

$$t^* = \min \left\{ t > \Delta : \max_{s \in [\Delta, t) \cap \mathcal{S}(t)} \left\{ \Delta \left(\frac{t-s}{ts} \right)^{1/2} \kappa - 2b_t \right\} \geq 0 \right\}.$$

Step 4. Write for convenience $m = t - \Delta$, so that $d \leq m$. Recalling that $b_t = 2^{3/2} \sigma \log^{1/2}(t/\alpha)$, we seek the smallest integer m such that

$$\max_{s \in [\Delta, m+\Delta) \cap \mathcal{S}(m+\Delta)} \left[\Delta \kappa \left\{ \frac{m+\Delta-s}{(m+\Delta)s} \right\}^{1/2} - 2^{5/2} \sigma \log^{1/2}\{(m+\Delta)/\alpha\} \right] > 0,$$

which is equivalent to

$$\max_{s \in [\Delta, m+\Delta) \cap \mathcal{S}(m+\Delta)} \left[\Delta^2 \kappa^2 - 32\sigma^2 \frac{s(m+\Delta)}{m+\Delta-s} \log\{(m+\Delta)/\alpha\} \right] > 0.$$

In turn, the above task corresponds to that of computing the smallest integer m such that

$$\Delta^2 \kappa^2 > \min_{s \in [\Delta, m+\Delta) \cap \mathcal{S}(m+\Delta)} \left[32\sigma^2 \frac{s(m+\Delta)}{m+\Delta-s} \log\{(m+\Delta)/\alpha\} \right].$$

Since $[\Delta, m+\Delta) \cap \mathcal{S}(m+\Delta) = \{m+\Delta-2^j, j = 0, \dots, \lfloor \log(m)/\log(2) \rfloor\}$, in the following we consider two cases.

Case 1. If $\log(m)/\log(2)$ is an integer, then it follows from identical arguments in the proof of [Theorem 2.2](#) that

$$\begin{aligned} \Delta^2 \kappa^2 &> \min_{s \in [\Delta, m+\Delta) \cap \mathcal{S}(t)} \left[32\sigma^2 \frac{s(m+\Delta)}{m+\Delta-s} \log\{(m+\Delta)/\alpha\} \right] \\ &= 32\sigma^2 \frac{\Delta(m+\Delta)}{m} \log\{(m+\Delta)/\alpha\}. \end{aligned}$$

Case 2. If $\log(m)/\log(2)$ is not an integer, then $\lfloor \log(m)/\log(2) \rfloor < \log(m)/\log(2)$. This means

$$\begin{aligned} \Delta^2 \kappa^2 &> \min_{s \in [\Delta, m+\Delta) \cap \mathcal{S}(t)} \left[32\sigma^2 \frac{s(m+\Delta)}{m+\Delta-s} \log\{(m+\Delta)/\alpha\} \right] \\ &> 32\sigma^2 \frac{\Delta(m+\Delta)}{m} \log\{(m+\Delta)/\alpha\}. \end{aligned}$$

Both cases lead to finding the smallest integer m such that

$$m \left[\frac{\Delta \kappa^2}{32\sigma^2} - \log\{(m+\Delta)/\alpha\} \right] > \Delta \log\{(m+\Delta)/\alpha\}.$$

It follows from the identical arguments in the proof of [Theorem 2.2](#) that we complete the proof. $\square$

Proof of Theorem 3.3. This proof consists of four steps. In **Step 1** we define a large probability event, where the rest of the proof is conducted. In **Step 2** we conclude the claims on type I error controls. In **Step 3** we derive the detection delay upper bounds.

Step 1. To prove the theorem, it is sufficient (in fact, equivalent) to show that the claim holds true simultaneously over all deterministic stopping times. Toward that end, define the event

$$\mathcal{E} = \left\{ \forall e, s, t \in \mathbb{N}, e \geq 0, s > e, t > s : \left| \left\{ \frac{t-s}{(s-e)(t-e)} \right\}^{1/2} \sum_{l=e+1}^s (X_l - f_l) - \left\{ \frac{s-e}{(t-s)(t-e)} \right\}^{1/2} \sum_{l=s+1}^t (X_l - f_l) \right| < b_t \right\}.$$

Then, under $\mathcal{E}$, the fluctuations of the CUSUM process are controlled uniformly at any time t , including any stopping (in fact, more generally, random) time T . It follows from [Lemma A.2](#) that

$$\mathbb{P}\{\mathcal{E}\} > 1 - \alpha.$$

On the event $\mathcal{E}$, for any $0 \leq e < s < t$, it holds that

$$\left| \widehat{D}_{e,s,t} - D_{e,s,t} \right| < b_t,$$

which implies that

$$D_{e,s,t} + b_{e,t} > \widehat{D}_{e,s,t} > D_{e,s,t} - b_{e,t}. \quad (\text{B.7})$$

In addition, due to [Assumption 3.2](#), we have that

$$C_d \frac{\sigma^2 \log(\Delta/\alpha)}{\kappa^2} \leq \Delta/4. \quad (\text{B.8})$$

Then it suffices to show that

- i. for any refresh starting point of the algorithm e and any interval $(e, t]$ not containing any true change points, on the event $\mathcal{E}$, there is no detected change point;
- ii. on the event $\mathcal{E}$, we can detect η_k with delay upper bounded by $C_d \sigma^2 \log(\Delta/\alpha) \kappa_k^{-2}$.

Step 2. As for (i), it holds automatically due to the definition of the event $\mathcal{E}$. The claim (i) leads to that $\widehat{t}_k > \eta_k$.

Step 3. As for (ii), we prove by induction. When $k=0$, we have $\widehat{t}_k = \eta_k = 0$, then $\eta_1 - \widehat{t}_0 \geq \Delta \geq 3\Delta/4$. It follows from identical arguments in the proof of [Theorem 2.2](#) that

$$d_1 = \widehat{t}_1 - \eta_1 \leq C_d \frac{\sigma^2 \log(\Delta/\alpha)}{\kappa_1^2}.$$

Due to (B.8), we have that $\eta_2 - \widehat{t}_1 \geq 3\Delta/4$, then due to [Algorithm 4](#), the procedure restarts by setting $e = \widehat{t}_1$. For a general $k \geq 1$, if $\eta_k - \widehat{t}_{k-1} \geq 3\Delta/4$, then it follows from identical arguments in the proof of [Theorem 2.2](#) that

$$d_k = \widehat{t}_k - \eta_k \leq C_d \frac{\sigma^2 \log(\Delta/\alpha)}{\kappa_k^2},$$

which completes the proof. □

Proof of Proposition 4.1. Throughout the proof we will assume for simplicity that $\frac{\sigma^2}{2\kappa^2} \log(\frac{1}{\alpha})$ is an integer. This proof consists of four steps. In **Step 1** we apply the change of variable argument and deal with the event when the sum of Z_i 's are upper bounded. In **Step 2**, we deal with the harder event when the sum of Z_i 's are lower bounded. The two events are combined in **Step 3** and the proof is concluded in Step 4.

Step 1. For any n and ν , let P^n be the restrictions of a distribution P to $\mathcal{F}_n$, i.e., the σ -field generated by the observations $\{X_i\}_{i=1}^n$. For any $\nu \geq 1$ and $n \geq \nu$, we have that

$$\frac{dP_{\kappa, \sigma, \nu}^n}{dP_{\kappa, \sigma, \infty}^n} = \exp \left(\sum_{i=\nu+1}^n Z_i \right),$$

where $P_{\kappa, \sigma, \infty}$ indicates the joint distribution under which there is no change point and

$$Z_i = \frac{\mu_2 - \mu_1}{\sigma^2} \left(X_i - \frac{\mu_1 + \mu_2}{2} \right).$$

For any $\nu \geq 1$, define the event

$$\mathcal{E}_\nu = \left\{ \nu < T < \nu + \frac{\sigma^2}{2\kappa^2} \log \left(\frac{1}{\alpha} \right), \sum_{i=\nu+1}^T Z_i < \frac{3}{4} \log \left(\frac{\nu}{\alpha} \right) \right\}.$$

Then we have

$$\begin{aligned} \mathbb{P}_{\kappa, \sigma, \nu}(\mathcal{E}_\nu) &= \int_{\mathcal{E}_\nu} \exp \left(\sum_{i=\nu+1}^T Z_i \right) dP_{\kappa, \sigma, \infty} \leq \exp \{ (3/4) \log (\nu/\alpha) \} \mathbb{P}_{\kappa, \sigma, \infty}(\mathcal{E}_\nu) \\ &\leq \exp \{ (3/4) \log (\nu/\alpha) \} \mathbb{P}_{\kappa, \sigma, \infty} \left\{ \nu < T < \nu + \frac{\sigma^2}{2\kappa^2} \log \left(\frac{\nu}{\alpha} \right) \right\} \\ &\leq \Delta^{3/4} \alpha^{-3/4} \mathbb{P}_{\kappa, \sigma, \infty} \left\{ \nu < T < \nu + \frac{\sigma^2}{2\kappa^2} \log \left(\frac{\nu}{\alpha} \right) \right\}, \end{aligned} \quad (\text{B.9})$$

where the first two inequalities follow from the definition of $\mathcal{E}_\nu$.

Step 2. For any $\nu \geq 1$ and $T \in \mathcal{D}(\alpha)$, since $\{T \geq \nu\} \in \mathcal{F}_{\nu-1}$, we have that

$$\begin{aligned} &\mathbb{P}_{\kappa, \sigma, \nu} \left\{ \nu < T < \nu + \frac{\sigma^2}{2\kappa^2} \log \left(\frac{\nu}{\alpha} \right), \sum_{i=\nu+1}^T Z_i \geq (3/4) \log (\nu/\alpha) \mid T > \nu \right\} \\ &\leq \text{esssup } \mathbb{P}_{\kappa, \sigma, \nu} \left\{ \max_{1 \leq t \leq \frac{\sigma^2}{2\kappa^2} \log (\frac{\nu}{\alpha}) - 1} \sum_{i=\nu+1}^{\nu+t} Z_i \geq (3/4) \log (\nu/\alpha) \mid X_1, \dots, X_\nu \right\} \\ &\leq \text{esssup } \mathbb{P}_{\kappa, \sigma, \nu} \left[\max_{1 \leq t \leq \frac{\sigma^2}{2\kappa^2} \log (\frac{\nu}{\alpha}) - 1} \sum_{i=\nu+1}^{\nu+t} \{Z_i - \kappa^2/(2\sigma^2)\} \geq (3/4) \log (\nu/\alpha) \right. \\ &\quad \left. - \frac{\sigma^2}{2\kappa^2} \log \left(\frac{\nu}{\alpha} \right) \frac{\kappa^2}{2\sigma^2} \mid X_1, \dots, X_\nu \right] \\ &\leq \text{esssup } \mathbb{P}_{\kappa, \sigma, \nu} \left[\max_{1 \leq t \leq \frac{\sigma^2}{2\kappa^2} \log (\frac{\nu}{\alpha}) - 1} \sum_{i=\nu+1}^{\nu+t} \{Z_i - \kappa^2/(2\sigma^2)\} \geq (1/2) \log (\nu/\alpha) \mid X_1, \dots, X_\nu \right] \\ &\leq \frac{\sigma^2}{2\kappa^2} \log \left(\frac{\nu}{\alpha} \right) \exp \left\{ - \frac{(1/2) \log^2 (\nu/\alpha)}{\frac{\sigma^2}{2\kappa^2} \log (\frac{\nu}{\alpha}) \frac{\kappa^2}{\sigma^2}} \right\} = \frac{\sigma^2}{2\kappa^2} \log \left(\frac{\nu}{\alpha} \right) \exp \{ - \log (\nu/\alpha) \} \leq \alpha^{1/4}, \end{aligned}$$

where the fourth inequality follows from Hoeffding's inequality and a union bound argument, and the last inequality holds for all small enough α such that

$$2\kappa^2 \nu \sigma^{-2} \geq \log (\nu/\alpha) \alpha^{3/4}.$$

We therefore let

$$\mathcal{P} = \{P_{\kappa, \sigma, \nu} : 2\kappa^2 \nu \sigma^{-2} \geq \log (\nu/\alpha) \alpha^{3/4}\}.$$

Since the upper bound is independent of ν , κ and σ , it holds that

$$\sup_{P_{\kappa, \sigma, \nu} \in \mathcal{P}} \mathbb{P}_{\kappa, \sigma, \nu} \left\{ \nu < T < \nu + \frac{\sigma^2}{2\kappa^2} \log \left(\frac{\nu}{\alpha} \right), \sum_{i=\nu+1}^T Z_i \geq (3/4) \log(\nu/\alpha) \mid T \geq \nu \right\} \leq \alpha^{1/4},$$

which leads to

$$\sup_{P_{\kappa, \sigma, \nu} \in \mathcal{P}} \mathbb{P}_{\kappa, \sigma, \nu} \left\{ \nu < T < \nu + \frac{\sigma^2}{2\kappa^2} \log \left(\frac{\nu}{\alpha} \right), \sum_{i=\nu+1}^T Z_i \geq (3/4) \log(\nu/\alpha) \right\} \leq \alpha^{1/4} \quad (\text{B.10})$$

Combining (B.9) and (B.10), we have

$$\begin{aligned} & \sup_{P_{\kappa, \sigma, \nu} \in \mathcal{P}} \mathbb{P}_{\kappa, \sigma, \nu} \left\{ \nu < T < \nu + \frac{\sigma^2}{2\kappa^2} \log \left(\frac{\nu}{\alpha} \right) \right\} \\ & \leq \Delta^{3/4} \alpha^{-3/4} \mathbb{P}_{\kappa, \sigma, \infty} \left\{ \nu < T < \nu + \frac{\sigma^2}{2\kappa^2} \log \left(\frac{\nu}{\alpha} \right) \right\} + \alpha^{1/4}. \end{aligned} \quad (\text{B.11})$$

Step 3. We now have, for any change point time Δ ,

$$\begin{aligned} & \inf_{\widehat{t} \in \mathcal{D}(\alpha)} \sup_{P_{\kappa, \sigma, \Delta} \in \mathcal{P}} \mathbb{P}_{\kappa, \sigma, \Delta} \left\{ \widehat{t} - \Delta \geq \frac{\sigma^2}{2\kappa^2} \log(\Delta/\alpha) \right\} \\ & = \inf_{\widehat{t} \in \mathcal{D}(\alpha)} \sup_{P_{\kappa, \sigma, \Delta} \in \mathcal{P}} \left[\mathbb{P}_{\kappa, \sigma, \Delta} \{ \widehat{t} > \Delta \} - \mathbb{P}_{\kappa, \sigma, \Delta} \left\{ \Delta < \widehat{t} < \Delta + \frac{\sigma^2}{2\kappa^2} \log \left(\frac{\Delta}{\alpha} \right) \right\} \right] \\ & \geq \inf_{\widehat{t} \in \mathcal{D}(\alpha)} \sup_{P_{\kappa, \sigma, \Delta} \in \mathcal{P}} \left[1 - \alpha - \alpha^{1/4} - \Delta^{3/4} \alpha^{-3/4} \mathbb{P}_{\kappa, \sigma, \infty} \left\{ \Delta < \widehat{t} < \Delta + \frac{\sigma^2}{2\kappa^2} \log \left(\frac{\Delta}{\alpha} \right) \right\} \right] \\ & \geq c_1 - \sup_{\widehat{t} \in \mathcal{D}(\alpha)} \inf_{P_{\kappa, \sigma, \Delta} \in \mathcal{P}} \Delta^{3/4} \alpha^{-3/4} \mathbb{P}_{\kappa, \sigma, \infty} \left\{ \Delta < \widehat{t} < \Delta + \frac{\sigma^2}{2\kappa^2} \log \left(\frac{\Delta}{\alpha} \right) \right\} \geq c. \end{aligned}$$

where the first inequality is due to the definition of the class of $\mathcal{D}(\alpha)$ of stopping times (which in particular implies that $\mathbb{P}_{\kappa, \sigma, \Delta} \{ \widehat{t} \leq \Delta \} = \mathbb{P}_{\infty}(\widehat{t} \leq \Delta) \leq \mathbb{P}_{\infty}(\widehat{t} < \infty) \leq \alpha$); the second inequality holds for an absolute constant $c_1 > 0$, provided that

$$1 - c_1 > \alpha + \alpha^{1/4},$$

and the final inequality holds provided that

$$\sup_{\widehat{t} \in \mathcal{D}(\alpha)} \inf_{P_{\kappa, \sigma, \Delta} \in \mathcal{P}} \mathbb{P}_{\kappa, \sigma, \infty} \left\{ \Delta < \widehat{t} < \Delta + \frac{\sigma^2}{2\kappa^2} \log \left(\frac{\Delta}{\alpha} \right) \right\} \leq \alpha/\Delta \quad (\text{B.12})$$

and that $\alpha^{1/4} \Delta^{-1/4} < c_1 - c$, with $c > 0$ being an absolute constant.

Step 4. It now suffices to show that (B.12) holds. We now prove by contradiction and assume that

$$\sup_{\widehat{t} \in \mathcal{D}(\alpha)} \inf_{P_{\kappa, \sigma, \Delta} \in \mathcal{P}} \mathbb{P}_{\kappa, \sigma, \infty} \left\{ \Delta < \widehat{t} < \Delta + \frac{\sigma^2}{2\kappa^2} \log \left(\frac{\Delta}{\alpha} \right) \right\} > \alpha/\Delta,$$

which implies that there exists $\widehat{t} \in \mathcal{D}(\alpha)$ such that

$$\inf_{\Delta \geq 1} \mathbb{P}_{\kappa, \sigma, \infty} \left\{ \Delta < \widehat{t} < \Delta + \frac{\sigma^2}{2\kappa^2} \log \left(\frac{\Delta}{\alpha} \right) \right\} > \alpha/\Delta,$$

which contradicts with the definition of $\mathcal{D}(\alpha)$ that

$$\mathbb{P}_{\kappa, \sigma, \infty} \{ \widehat{t} < \infty \} < \alpha.$$

We therefore conclude the proof. □

Table C1. Numerical results.

α	σ	Method	PFA	DD	α	σ	Method	PFA	DD
0.05	0.1	M	0.02	0.00	0.1	0.1	M	0.00	0.00
0.05	0.1	M2	0.08	0.00	0.1	0.1	M2	0.10	0.00
0.05	0.1	M3	0.10	0.00	0.1	0.1	M3	0.00	0.00
0.05	0.5	M	0.00	5.33	0.1	0.05	M	0.02	4.92
0.05	0.5	M2	0.04	3.78	0.1	0.0	M2	0.11	3.28
0.05	0.5	M3	0.00	6.45	0.1	0.05	M3	0.00	5.66
0.05	0.8	M	0.02	15.59	0.1	0.8	M	0.10	14.24
0.05	0.8	M2	0.10	12.92	0.1	0.8	M2	11.84	8.78
0.05	0.8	M3	0.00	22.96	0.1	0.8	M3	20.38	8.78
0.05	1.2	M	0.01	80.09	0.1	1.2	M	0.03	66.52
0.05	1.2	M2	0.04	73.6	0.1	1.2	M2	0.11	51.21
0.05	1.2	M3	0.00	75.4	0.1	1.2	M3	0.00	51.36

For both metrics the numbers are obtained by averaging over 100 Monte Carlo simulations. PFA, probability of false alarm; DD, detection delay.

APPENDIX C. ADDITIONAL EXPERIMENTS: VARIATIONS OF MAILLARD (2019)

Following an anonymous referee's comment, we conclude with a simple experiment consisting of variations of the method from Maillard (2019). Recall the notation in Maillard (2019). For $t_2 > 1$, let

$$\mu_{(t_1+1):t_2} = \frac{1}{t_2 - t_1} \sum_{s=t_1+1}^{t_2} X_s.$$

Maillard (2019) proposes to raise an alarm at $\hat{t}$ with

$$\hat{t} = \inf\{t > 1 : \exists s \in [1, t) \text{ s.t. } |\mu_{0:s} - \mu_{(s+1):t}| \geq b_{s,t}\},$$

where

$$b_{s,t} = \sigma \sqrt{2 \left(\frac{1}{s} + \frac{1}{t-s} \right) \left(1 + \frac{1}{t} \right) \log \left(\frac{2(t-1)\sqrt{t-1}}{\alpha} \right)}.$$

We refer to this procedure as M in the prequel. In this section, we consider two variations thereof.

The first variation, referred as M2, consists of replacing $b_{s,t}$ with

$$1.1\sigma \sqrt{2 \left(\frac{1}{s} + \frac{1}{t-s} \right) \left(1 + \frac{1}{t} \right) \log \left(\frac{2(t-1)\sqrt{t-1}}{\alpha} \right)}.$$

The second variation, namely M3, replaces $b_{s,t}$ with

$$0.9\sigma \sqrt{2 \left(\frac{1}{s} + \frac{1}{t-s} \right) \left(1 + \frac{1}{t} \right) \log \left(\frac{2(t-1)\sqrt{t-1}}{\alpha} \right)}.$$

For the setting considered in Section 5, the results collected in Table C1 show that M2 and M3 do not lead to a desired PFA similar to the prespecified α . Thus, at least to us, it is unclear how to improve the detection method from Maillard (2019). In particular, small changes in the criterion from Maillard (2019) lead to significant changes in the PFA but no substantial improvement in the detection delay.

REFERENCES

Aue, Alexander, and Lajos Horváth. 2004. "Delay Time in Sequential Detection of Change." *Statistics & Probability Letters* 67 (3): 221–231. <https://doi.org/10.1016/j.spl.2004.01.002>

- Aue, Alexander, Lajos Horváth, and Matthew L. Reimherr. 2009. "Delay Times of Sequential Procedures for Multiple Time Series Regression Models." *Journal of Econometrics* 149 (2): 174–190. <https://doi.org/10.1016/j.jeconom.2008.12.018>
- Chen, Guangliang, Stefan Atev, and Gilad Lerman. 2009. "Kernel Spectral Curvature Clustering (KSCC)." In 2009 IEEE 12th International Conference on Computer Vision Workshops, ICCV Workshops, Kyoto, Japan, 27 September–4 October 2009, 765–772. IEEE.
- Chen, Hao. 2019. "Sequential Change-Point Detection Based on Nearest Neighbors." *The Annals of Statistics* 47 (3): 1381–1407. <https://doi.org/10.1214/18-AOS1718>
- Chen, Hao, and Lynna Chu. 2019. "gStream: Graph-Based Sequential Change-Point Detection for Streaming Data." R package version 0.2.0. <https://cran.r-project.org/web/packages/gStream/index.html>
- Chen, Yudong, Tengyao Wang, and Richard J. Samworth. 2020. "High-Dimensional, Multiscale Online Changepoint Detection." arXiv preprint arXiv: 2003.03668.
- Chu, Chia-Shang James, Maxwell Stinchcombe, and Halbert White. 1996. "Monitoring Structural Change." *Econometrica* 64 (5): 1045. <https://doi.org/10.2307/2171955>
- Desobry, Frédéric, Manuel Davy, and Christian Doncarli. 2005. "An Online Kernel Change Detection Algorithm." *IEEE Transactions on Signal Processing* 53 (8): 2961–2974. <https://doi.org/10.1109/TSP.2005.851098>
- Dette, Holger, and Josua Gösmann. 2020. "A Likelihood Ratio Approach to Sequential Change Point Detection for a General Class of Parameters." *Journal of the American Statistical Association* 115 (531): 1361–1377. <https://doi.org/10.1080/01621459.2019.1630562>
- Fearnhead, Paul, and Zhen Liu. 2007. "On-Line Inference for Multiple Changepoint Problems." *Journal of the Royal Statistical Society Series B: Statistical Methodology* 69 (4): 589–605. <https://doi.org/10.1111/j.1467-9868.2007.00601.x>
- Gösmann, Josua, Tobias Kley, and Holger Dette. 2019. "A New Approach for Open-End Sequential Change Point Monitoring." arXiv preprint arXiv: 1906.03225.
- He, Xi, Yao Xie, Sin-Mei Wu, and Fan-Chi Lin. 2018. "Sequential Graph Scanning Statistic for Change-Point Detection." In 2018 52nd Asilomar Conference on Signals, Systems, and Computers, Pacific Grove, CA, USA, 1317–1321. IEEE. <https://doi.org/10.1109/ACSSC.2018.8645505>
- Hlávka, Zdenek, Marie Husková, Claudia Kirch, and Simos G. Meintanis. 2016. "Bootstrap Procedures for Online Monitoring of Changes in Autoregressive Models." *Communications in Statistics—Simulation and Computation* 45 (7): 2471–2490. <https://doi.org/10.1080/03610918.2014.904346>
- Howard, Steven R., Aaditya Ramdas, Jon McAuliffe, and Jasjeet Sekhon. 2018. "Uniform, Nonparametric, Non-Asymptotic Confidence Sequences." arXiv preprint arXiv: 1810.08240.
- Husková, Marie, and Claudia Kirch. 2012. "Bootstrapping Sequential Change-Point Tests for Linear Regression." *Metrika* 75 (5): 673–708. <https://doi.org/10.1007/s00184-011-0347-7>
- Husková, Marie, Claudia Kirch, and Simos G. Meintanis. 2010. "Fourier Methods for Sequential Change Point Analysis in Autoregressive Models." In *Proceedings of COMPSTAT'2010*, 501–508. Limassol, Cyprus: Springer.
- Keshavarz, Hossein, George Michailidis, and Yves Atchade. 2018. "Sequential Change-Point Detection in High-Dimensional Gaussian Graphical Models." arXiv preprint arXiv: 1806.07870.
- Kirch, Claudia. 2008. "Bootstrapping Sequential Change-Point Tests." *Sequential Analysis* 27 (3): 330–349. <https://doi.org/10.1080/07474940802241082>
- Kirch, Claudia, and Silke Weber. 2018. "Modified Sequential Change Point Procedures Based on Estimating Functions." *Electronic Journal of Statistics* 12 (1): 1579–1613. <https://doi.org/10.1214/18-EJS1431>
- Kurt, Mehmet Necip, Yasin Yilmaz, and Xiaodong Wang. 2018. "Real-Time Nonparametric Anomaly Detection in High-Dimensional Settings." arXiv preprint arXiv: 1809.05250.
- Lai, Tze Leung. 1981. "Asymptotic Optimality of Invariant Sequential Probability Ratio Tests." *The Annals of Statistics* 9 (2): 318–333. <https://doi.org/10.1214/aos/1176345398>
- Lai, Tze Leung. 1995. "Sequential Changepoint Detection in Quality Control and Dynamical Systems." *Journal of the Royal Statistical Society: Series B (Methodological)* 57 (4): 613–644. <https://doi.org/10.1111/j.2517-6161.1995.tb02052.x>

- Lai, Tze Leung. 1998. "Information Bounds and Quick Detection of Parameter Changes in Stochastic Systems." *IEEE Transactions on Information Theory* 44 (7): 2917–2929.
- Lai, Tze Leung. 2001. "Sequential Analysis: some Classical Problems and New Challenges." *Statistica Sinica* 11 (2): 303–351.
- Lorden, Gary. 1971. "Procedures for Reacting to a Change in Distribution." *The Annals of Mathematical Statistics* 42 (6): 1897–1908. <https://doi.org/10.1214/aoms/1177693055>
- Maillard, Odalric-Ambrym. 2019. "Sequential Change-Point Detection: Laplace Concentration of Scan Statistics and Non-Asymptotic Delay Bounds." In *Algorithmic Learning Theory*, 610–632. PMLR.
- Mei, Yajun. 2010. "Efficient Scalable Schemes for Monitoring a Large Number of Data Streams." *Biometrika* 97 (2): 419–433. <https://doi.org/10.1093/biomet/asq010>
- Moustakides, George V. 1986. "Optimal Stopping Times for Detecting Changes in Distributions." *The Annals of Statistics* 14 (4): 1379–1387. <https://doi.org/10.1214/aos/1176350164>
- Namoano, Bernadin, Andrew Starr, Christos Emmanouilidis, and Ruiz Carcel Cristobal. 2019. "Online Change Detection Techniques in Time Series: An Overview." In 2019 IEEE International Conference on Prognostics and Health Management (ICPHM), San Francisco, CA, USA, 1–10. IEEE. <https://doi.org/10.1109/ICPHM.2019.8819394>
- Page, Ewan S. 1954. "Continuous Inspection Schemes." *Biometrika* 41 (1-2): 100–115. <https://doi.org/10.1093/biomet/41.1-2.100>
- Ritov, Ya'acov. 1990. "Decision Theoretic Optimality of the CUSUM Procedure." *The Annals of Statistics* 18 (3): 1464–1469. <https://doi.org/10.1214/aos/1176347761>
- Siegmund, David. 2013. *Sequential Analysis: tests and Confidence Intervals*. New York: Springer Science & Business Media.
- Tartakovsky, Alexander, Igor Nikiforov, and Michele Basseville. 2014. *Sequential Analysis: Hypothesis Testing and Changepoint Detection*. Boca Raton, FL: Chapman and Hall/CRC.
- Wainwright, Martin J. 2019. *High-Dimensional Statistics: A Non-Asymptotic Viewpoint*. Vol. 48. Cambridge, UK: Cambridge University Press.
- Wald, A. 1945. "Sequential Tests of Statistical Hypotheses." *The Annals of Mathematical Statistics* 16 (2): 117–186. <https://doi.org/10.1214/aoms/1177731118>
- Wang, Daren, Yi Yu, and Alessandro Rinaldo. 2020. "Univariate Mean Change Point Detection: Penalization, Cusum and Optimality." *Electronic Journal of Statistics* 14 (1): 1917–1961. <https://doi.org/10.1214/20-EJS1710>