

Window-Limited CUSUM for Sequential Change Detection

Liyan Xie¹, George V. Moustakides², *Life Senior Member, IEEE*, and Yao Xie³, *Member, IEEE*

Abstract—We study the parametric online changepoint detection problem, where the underlying distribution of the streaming data changes from a known distribution to an alternative that is of a known parametric form but with unknown parameters. We propose a joint detection/estimation scheme, which we call Window-Limited CUSUM, that combines the cumulative sum (CUSUM) test with a sliding window-based consistent estimate of the post-change parameters. We characterize the optimal choice of the window size and show that the Window-Limited CUSUM enjoys first-order asymptotic optimality as average run length approaches infinity under the optimal choice of window length. Compared to existing schemes with similar asymptotic optimality properties, our test can be much faster computed because it can recursively update the CUSUM statistic by employing the estimate of the post-change parameters. A parallel variant is also proposed that facilitates the practical implementation of the test. Numerical simulations corroborate our theoretical findings.

Index Terms—Cumulative sum (CUSUM) test, sequential change detection, average run length, asymptotic optimality.

I. INTRODUCTION

ONLINE changepoint detection is a fundamental problem in statistics and signal processing [1], [2], [3], [4] which finds applications in a plethora of practical problems in diverse fields. The most common version of the problem consists of a sequence of observations sampled independently. There is also a changepoint such that the underlying distribution changes from one distribution to an alternative. This problem is of major importance in many applications, such as seismic signal processing [5], industrial quality control [6], dynamical systems monitoring [7], structural health control [8], event detection [9], anomaly detection [10], detection of attacks [11], etc. The goal of online changepoint detection is to detect

the occurrence of the change in statistical behavior with a minimal delay while controlling the false alarm rate. The suitable tradeoff between detection delay and false alarm rate, as in all detection problems, is of essential importance for the proper mathematical formulation of the problem.

Classical formulations assume complete knowledge of the pre- and post-change underlying distributions with the cumulative sum (CUSUM) test being the most popular means for the corresponding detection [12]. The CUSUM scheme properly updates the log-likelihood ratio between the post- and pre-change densities of the available data to form the corresponding test statistic. CUSUM is known to be theoretically optimum in the sense that it enjoys minimum detection delay under a fixed false alarm rate constraint [13]. Also, it is computationally simple because an updating formula exists for the computation of the CUSUM statistics, which requires only the current data sample (and the previous value of the CUSUM statistic).

In many real world applications, the post-change distribution is typically not precisely known since it represents a switching to an anomalous state. In this case, a Generalized Likelihood Ratio (GLR) test version of CUSUM has been developed, which applies the GLR method to select the unknown post-change distributions [14] and form the corresponding test statistic. Unfortunately, this original GLR version turns out to be computationally demanding because it requires computations per sample, which increase linearly in time without limit. The main reason for this disadvantage is that the test statistic must be recomputed using GLR for each possible changepoint location with every new observation. A remedy also proposed in [14] is the window-limited GLR test where the recomputation of the test statistic is limited to changepoint locations within a window of fixed length starting from the current time instant. The resulting scheme has been shown to enjoy asymptotic optimality with proper selection of the window size. Even though computations are now limited since they are of the order of the selected window length, they still tend to be considerable because for each time instant we need to recompute the GLR statistic for *each* position within the window.

In this article, we develop an alternative approach for solving the problem of interest which we call Window-Limited CUSUM (WLCUSUM). It consists in adopting a window-based estimate of the unknown post-change parameters and, unlike the existing window-limited GLR, we use the estimate in the updating formula of the classical

Manuscript received 6 November 2022; revised 11 March 2023; accepted 28 April 2023. Date of publication 19 May 2023; date of current version 18 August 2023. The work of Yao Xie was supported in part by an NSF CAREER under Grant CCF-1650913; and in part by NSF under Grant DMS-2134037, Grant CMMI-2015787, Grant CMMI-2112533, Grant DMS-1938106, and Grant DMS-1830210. (*Corresponding author: Liyan Xie.*)

Liyan Xie is with the School of Data Science, The Chinese University of Hong Kong, Shenzhen 518172, China (e-mail: xielian@cuhk.edu.cn).

George V. Moustakides is with the Electrical and Computer Engineering Department, University of Patras, 26500 Patras, Greece (e-mail: moustaki@upatras.gr).

Yao Xie is with the H. Milton Stewart School of Industrial and Systems Engineering, Georgia Institute of Technology, Atlanta, GA 30332 USA (e-mail: yao.xie@isye.gatech.edu).

Communicated by L. Lai, Associate Editor for Signal Processing and Source Coding.

Color versions of one or more figures in this article are available at <https://doi.org/10.1109/TIT.2023.3274646>.

Digital Object Identifier 10.1109/TIT.2023.3274646

CUSUM statistic. This updating mechanism is far more efficient than its window-limited GLR counterpart and by proper selection of the window size we can also guarantee asymptotic optimality. In detail, we only require the sample size $w \rightarrow \infty$ as $\gamma \rightarrow \infty$ with $w = o(\log \gamma)$ while the window-limited GLR require $w = \Theta(\log \gamma)$ where γ is the average run length requirement (see Remark 8 for details). We would like to emphasize that the problem we consider is not joint detection and estimation as in [15] and [16], where the two tasks are regarded as equally important. Here, we are primarily interested in detection, with estimation being an auxiliary action that contributes towards our detection goal. For this reason we only require the estimator to be consistent without insisting on any explicit form.

Compared with existing CUSUM-like procedures employing estimates of post-change parameters, the proposed WLCUSUM method applies to a far wider range of parametric distributions and not only the exponential family which is mostly the case with the available approaches. Our main contributions in this work include: (i) Proof of asymptotic optimality of the proposed WLCUSUM procedure under Lorden's *worst-case* detection delay [17]. To achieve this goal, we had to develop new upper bounds for overshoots over a constant threshold for sums of data that are w -dependent. (ii) Characterization of the *optimal* choice of the window size to guide practical implementations and offer the best possible performance for the proposed scheme. (iii) Development of an alternative parallel version of WLCUSUM capable of matching the performance of the optimal window without the need to explicitly specify it.

We must also mention that one of the main characteristics of WLCUSUM is its computational efficiency. The benchmark window-limited GLR [14] requires a window size that is at least as large as the detection delay, while in our detector the optimal window has a size that is significantly smaller. This difference in window size translates into an overall computational complexity which, in our scheme is at least an order of magnitude smaller than the corresponding complexity of the window-limited GLR.

The remainder of this paper is organized as follows. Section II reviews related work. Section III introduces the adopted formulation and presents details of the proposed detection procedure. Section V contains the theoretical analysis establishing the asymptotic optimality of the proposed procedure and the form of the optimal window size. Section VI presents a parallel implementation of the proposed WLCUSUM with varying window sizes which is particularly suited for practical implementation. Finally, Section VII presents examples that demonstrate the performance of the WLCUSUM procedure with comparisons to the corresponding GLR scheme. For better readability of the main results, most of our technical part is moved to the Appendix.

II. RELATED WORK

The study of online changepoint detection can be traced back to the early work of Page [12] and has been studied for several decades. Most articles consider the problem under independent observations but there are also extensions to more

complicated data models; see [1], [2], [3], [18], [19], [20] for thorough reviews in this field.

We distinguish two main tracks in sequential change detection: The first is the Bayesian approach, where a prior for the time of change is assumed to be available. The second is the minimax (non-Bayesian) formulation, where the changepoint is considered to be deterministic but unknown. Interestingly, both approaches can be put under the same mathematical framework [21], [22] and, depending on the data model, we can decide which formulation is most suitable to be adopted.

The first exact optimality result in sequential change detection can be found in [23] where the focus is on detecting a change in the drift of a Brownian motion. Following a Bayesian approach, the change time is modeled as an independent random variable that is exponentially distributed. For non-Bayesian approaches, the CUSUM test is perhaps the most popular change detection algorithm for the classical setup of detecting a change from a known nominal to a known alternative density. CUSUM, also known as the Page test [12], was first shown to be asymptotically optimum in [17] when observations are i.i.d. before and after the change. The exact optimality of the CUSUM test under the same data model was established in [13]. An alternative detection procedure introduced in [24], even though it was developed by adopting a minimax approach, presents very strong similarities to the optimum Bayesian test developed in [23]. This sequential detector, known as the Shiryaev-Roberts-Pollak (SRP) test, enjoys a very strong asymptotic optimality property which, unfortunately, was proven not to be exact [25]. We must also mention that in the classical version of the problem, which we discussed so far, the computation of the corresponding test statistic of the CUSUM and the SRP test is straightforward and can be implemented very efficiently.

When we consider parametric density families where the parameters are unknown, CUSUM and SRP are used as prototypes to develop variants which, at best, can enjoy asymptotic optimality. In particular, when the pre-change density is known while the post-change contains unknown parameters, there are two classes of tests in the literature that address this problem: (i) The generalized likelihood ratio (GLR) approach [14], where the detection statistic, at each time, is computed by substituting the unknown parameter with its maximum likelihood estimate using all potential post-change data until the current time; (ii) The mixture likelihood ratio procedure [3, Pages 418-423], where the detection statistic is a weighted average of the corresponding log-likelihood ratio by assuming a weight (prior distribution) on the post-change parameters. Although the corresponding tests, as mentioned, enjoy asymptotic optimality properties, a major drawback is that their computational complexity can be high because their detection statistic cannot be computed efficiently. In addition, the score statistics were also used to avoid estimation of the unknown parameters [26].

For the popular CUSUM test, many variants of the traditional version were proposed to improve the computational efficiency when the post-change density contains unknown parameters. To detect the change over a wide range of mean

shifts in quality control, the combined usage of CUSUM and Shewhart charts was employed in [27] and [28], and the simultaneous use of multiple CUSUM procedures with different drift values was suggested in [29], [30], [31], and [32]. Moreover, the case with finitely many post-change distributions was considered and the joint detection/isolation algorithm was proposed based on multiple hypotheses sequential probability ratio test [33], [34]. A different method known as Adaptive CUSUM was first proposed in [29] and studied further in [35], [36], [37], [38], [39], and [40]. The Adaptive CUSUM continuously adjusts its statistic in order to efficiently signal a one-step-ahead forecast in deviation from its target value. For example, such a procedure has been considered in [41], where the estimate is based on online algorithms such as stochastic gradient descent, and the performance metrics are related to the regret bounds of the online estimators. The simple exponentially weighted moving average (EWMA) estimate is the most common selection for the one-step-ahead forecast. Optimality properties of the Adaptive CUSUM were considered in [42] where the first-order asymptotic optimality was established for the univariate exponential family while extensions appear in [43]. Finally, a multi-stream Adaptive CUSUM test was proposed in [44] establishing asymptotic optimality for the case of Gaussian distributions.

III. PRELIMINARIES

A. Problem Setup

Suppose we have access to the multivariate data sequence $\{\xi_t\}$ with $\xi_t \in \mathbb{R}^k$, which is sampled sequentially. We assume that there are two probability density functions (pdf) $f_\infty(\cdot)$, $f_0(\cdot)$ and a deterministic time $\tau \in \{0, 1, 2, \dots\}$ such that

$$\xi_t \stackrel{\text{i.i.d.}}{\sim} \begin{cases} f_\infty(\xi), & t = 1, 2, \dots, \tau, \\ f_0(\xi, \theta), & t = \tau + 1, \tau + 2, \dots \end{cases} \quad (1)$$

In other words τ is a changepoint, where the observations are i.i.d. before and including τ following the pdf $f_\infty(\cdot)$, while for times after τ they are again i.i.d. following $f_0(\cdot, \theta)$, which is characterized by an unknown parameter vector $\theta \in \Theta \subseteq \mathbb{R}^K$ where Θ is a known subset of \mathbb{R}^K . If there is no constraint on θ then we simply set $\Theta = \mathbb{R}^K$. Throughout this paper we assume that the pre-change distribution $f_\infty(\cdot)$ does *not* belong to the post-change distributions with parameter set Θ .

We denote with P_∞, E_∞ the probability measure and the corresponding expectation when all samples follow the pre-change distribution (i.e., the change happens at ∞), P_0^θ, E_0^θ when all data are under the post-change density with parameter θ (the change happens at 0) and finally with $P_\tau^\theta, E_\tau^\theta$ the measure and expectation induced when the change happens at time τ . We also denote with $\xi_{t_1}^{t_2}$ ($t_2 \geq t_1$) the collection of data $\{\xi_{t_1}, \dots, \xi_{t_2}\}$.

We assume that $f_\infty(\cdot)$ is *known* since usually it can be estimated from historical data by density estimation [45] methods. We can also assume that $f_\infty(\cdot)$ is partially known by extending our results to incorporate the estimation error for $f_\infty(\cdot)$. However, we note that this estimation error can be negligible as long as the volume of historical data is

sufficiently large. For the post-change density, we assume that $f_0(\cdot, \theta)$ has a known form, but the parameter vector θ is unknown, without any prior distribution that can capture its statistical behavior. Our goal is to detect the changepoint τ as quickly as possible from streaming data when it occurs, under the constraint that the false alarm rate is properly controlled.

A sequential change detection test simply consists of a *stopping time* T which denotes the time we stop and declare that a change took place before time T . The stopping time is adapted to the filtration $\{\mathcal{F}_t\}$, $\mathcal{F}_t = \sigma\{\xi_1, \dots, \xi_t\}$ with \mathcal{F}_0 denoting the trivial sigma-algebra. This assures that only available data are employed when we decide whether to stop or not at each time t . Our intention is to use the classical CUSUM test for the change detection problem. Since CUSUM requires exact knowledge of the pre- and post-change densities, we will replace the unknown parameter vector θ with a proper estimate over available data. This estimate will be renewed with every new data sample. Before presenting the details of our scheme, let us first recall the CUSUM test and its corresponding optimality properties.

B. The CUSUM Test

Fix the post-change parameter vector θ and suppose it is *known*. This suggests that we are under the classical formulation, where we would like to detect a change from a known density $f_\infty(\cdot)$ to an alternative known density $f_0(\cdot, \theta)$. This problem can be solved optimally with the CUSUM test. The CUSUM statistic $\{S_t\}$ is defined as

$$S_t = \max_{0 \leq k < t} \sum_{j=k+1}^t \log \frac{f_0(\xi_j, \theta)}{f_\infty(\xi_j)},$$

which is essentially the (maximum) likelihood ratio statistics as detailed in [3, Section 8.2.3]. The above CUSUM statistic satisfies and is usually implemented through the following update

$$S_t = S_{t-1}^+ + \log \frac{f_0(\xi_t, \theta)}{f_\infty(\xi_t)}, \quad S_0 = 0, \quad (2)$$

where $x^+ = \max\{x, 0\}$. The update in (2) is applied every time a new sample becomes available. The corresponding CUSUM stopping time that signals the change is then defined as

$$T = \inf\{t > 0 : S_t \geq \nu\}, \quad (3)$$

where $\nu > 0$ is a constant threshold, the choice of which needs to balance the false alarm rate and the detection delay. The first time S_t hits or exceeds the threshold ν , we stop and declare that a change took place before t . CUSUM is known to solve *exactly* [13] the following challenging constrained optimization problem

$$\begin{aligned} & \inf_T \sup_{\tau \geq 0} \text{ess sup } E_\tau^\theta [T - \tau | \mathcal{F}_\tau, T > \tau], \\ & \text{subject to: } E_\infty[T] \geq \gamma > 1, \end{aligned} \quad (4)$$

when the threshold ν is selected to satisfy the false alarm constraint with equality, namely $E_\infty[T] = \gamma$. In other words, among all stopping times that have an average false alarm

period (also known as *average run length*, ARL) no smaller than γ , the CUSUM stopping time T has the smallest *worst-case average detection delay* (WADD). From (4) we observe that for each possible changepoint τ we consider the average detection delay conditioned on the *worst possible data* before (and including) time τ ; this is a particular type of delay measure proposed by Lorden [17]. It is well-known that CUSUM regarding this worst-case performance is an equalizer in the sense that $\text{ess sup } E_\tau^\theta[T - \tau | \mathcal{F}_\tau, T > \tau]$ is the same for all changepoints τ ; therefore, for the computation of the worst-case detection delay scenario, we can simply limit ourselves to $\tau = 0$ (i.e., $E_0^\theta[T]$).

The analysis in this article will be *asymptotic* (for large γ). For this reason, with the following lemma, we provide a convenient asymptotic formula for the CUSUM performance.

Lemma 1 (Performance of Exact CUSUM): For threshold $\nu = \log \gamma$ the CUSUM test satisfies

$$E_\infty[T] \geq \gamma, \quad E_0^\theta[T] = \frac{\log \gamma}{I_0} \left(1 + \Theta \left(\frac{1}{\log \gamma} \right) \right), \quad (5)$$

where $I_0 = E_0^\theta \left[\log \frac{f_0(\xi_1, \theta)}{f_\infty(\xi_1)} \right]$ is the Kullback-Leibler information number (divergence) of the post- and pre-change densities.

Proof: The performance of CUSUM, when expressed in asymptotic terms, is usually given in the form of $E_0^\theta[T] = \frac{\nu}{I_0} (1 + o(1))$ (see [2]). However, here we would like to be more explicit regarding the $o(1)$ term in order to be able to compare the case of known versus the case of estimated θ . The proof of this formula can be found in [46, Lemma 1]. \square

From Lemma 1 we conclude that if we know θ and apply the CUSUM test defined in (2), (3) with threshold $\nu = \log \gamma$ then the corresponding CUSUM stopping time T enjoys an asymptotic performance captured by (5). In fact no other stopping time T that satisfies the same false alarm constraint can have a limiting value (liminf) for the ratio $E_0^\theta[T] / \frac{\log \gamma}{I_0}$ that is smaller than 1 as $\gamma \rightarrow \infty$. This statement describes the optimality of CUSUM in first-order asymptotic terms as $\gamma \rightarrow \infty$.

IV. PROPOSED METHOD: WINDOW-LIMITED CUSUM TEST

In a realistic case θ is unknown and, as we mentioned, we may know instead a set Θ of possible values for θ . Of course $\Theta = \mathbb{R}^K$ if there is no restriction on θ . If θ is not exactly known then the CUSUM test cannot be applied in the form of (2), (3). For this reason, as in the literature, we propose to replace the unknown θ with a *consistent estimate*. Specifically we select a window of length $w \geq 1$ and define the Window-Limited CUSUM (WLCUSUM) test statistic $\{S_t\}$ for $t > w$ similarly to (2):

$$S_t = S_{t-1}^+ + \log \frac{f_0(\xi_t, \hat{\theta}_{t-1})}{f_\infty(\xi_t)}, \quad S_w = 0, \quad t = w+1, w+2, \dots, \quad (6)$$

where $\hat{\theta}_t \in \Theta$ is an estimate of θ . We are not going to adopt any specific estimator, we only constrain the estimate $\hat{\theta}_t$ to be based on the data $\{\xi_t, \dots, \xi_{t-w+1}\}$ and to be consistent.

An obvious possibility would be the Maximum Likelihood Estimator (MLE)

$$\hat{\theta}_t = \arg \max_{\theta \in \Theta} \sum_{i=0}^{w-1} \log f_0(\xi_{t-i}, \theta), \quad (7)$$

which, as we will see in our analysis, enjoys certain desirable optimality characteristics when employed in our proposed detection scheme. For the stopping time, similarly to CUSUM, we define

$$\mathcal{T} = \inf\{t > w : S_t \geq \nu\}, \quad (8)$$

with $\nu > 0$ is a constant threshold.

It is worth noting that the increment term $\{\log f_0(\xi_t, \hat{\theta}_{t-1}) / f_\infty(\xi_t)\}_{t \in \mathbb{Z}}$ in (6) is w -dependent (defined below); while the increment terms (log-likelihood ratios) in the exact CUSUM (2) are independent. Due to such w -dependency induced by the estimates which employ data from the past, obtaining the formulas for the detection performance is not a straightforward task.

Definition 1 (w -Dependence, [47]): A discrete-time stochastic process $(Y_n)_{n \in \mathbb{Z}}$ is w -dependent if for all k , the joint stochastic variables $(Y_n)_{n \leq k}$ are independent of the joint stochastic variables $(Y_n)_{n \geq k+w+1}$.

We must point out that several existing detection/estimation methods propose to recalculate the estimate $\hat{\theta}_t$ at each time t and perform a dual maximization. Such is, for example, the popular GLR approach proposed in [14]

$$\tilde{S}_t = \max_{0 \leq \tau < t} \sup_{\theta \in \Theta} \sum_{s=\tau+1}^t \log \frac{f_0(\xi_s, \theta)}{f_\infty(\xi_s)}.$$

Unfortunately, the above statistic does not possess any convenient updating formula similar to (2) and requires a number of operations per sample that increases linearly with time. To remedy this serious computational handicap, a *window-limited* version is commonly adopted where the search for the maximum over τ is performed within a window of fixed length w . Specifically, the following maximization replaces the previous one

$$\tilde{S}_t = \max_{t-w \leq \tau < t} \sup_{\theta \in \Theta} \sum_{s=\tau+1}^t \log \frac{f_0(\xi_s, \theta)}{f_\infty(\xi_s)}. \quad (9)$$

This clearly reduces the complexity to a fixed number of operations per time update but, as we discuss in Remark 8, it can still be quite demanding. In the following, we refer to (9) as the window-limited GLR approach.

Unlike the window-limited GLR, we propose to employ, as in [48], the classical CUSUM update in (2) where we simply replace the unknown parameter θ with a consistent estimate, thus preserving the computational efficiency of the original CUSUM. The reason we expect this idea to be successful is that when the data are under the post-change regime, $\hat{\theta}_{t-1}$ will be close to the true θ if w is sufficiently large, and the WLCUSUM statistic will exhibit a positive drift not differing significantly from the exact CUSUM drift. On the other hand, when the data follow the pre-change regime, we will show that the estimate will impose a negative drift on the WLCUSUM

statistic forcing the test to perform repeated restarts exactly similarly to the case of the exact CUSUM. These claims will be demonstrated through a rigorous analysis. Additionally, we will obtain asymptotic formulas for the average false alarm period and the worst-case average detection delay, which will allow us to establish the asymptotic optimality of our proposed detection scheme. Before starting our mathematical derivations, let us make some remarks and present our assumptions.

Remark 1: As can be seen from (6), we do not compute any test statistic during the first w samples since we accumulate these samples in order to obtain the first estimate $\hat{\theta}_w$. The test statistic is first computed at time $w + 1$ where we also employ the first estimate $\hat{\theta}_w$. Since we necessarily wait for w time instances, it is understood that whatever average delay we compute, it cannot be smaller than w . Asymptotically, this fact is not disturbing because we will assure in Section V that this initial waiting time w is *negligible* compared to the actual detection delay required to detect the change. Of course, not applying the test during the first w time instances is only for *analytical* purposes, as this corresponds to the worst-case average detection delay (as we prove in Lemma 4). In a practical implementation, we can start testing earlier, and our estimator at every time t can rely on the existing data without necessarily waiting until w samples become available. We discuss this point in more detail in Section VI when we introduce a computationally convenient variant of our scheme.

Remark 2: The estimate $\hat{\theta}_t$ that we employ in our test, as mentioned, is obtained by processing the samples $\{\xi_t, \dots, \xi_{t-w+1}\}$. This assures that ξ_t and $\hat{\theta}_{t-1}$ are *independent* and the same is true between $\hat{\theta}_t$ and \mathcal{F}_{t-w} (w -dependency). As we are going to see, these facts play an important role in our proofs.

Remark 3: We note that the name “window-limited CUSUM” has been used in the literature, e.g., [49] and [14]. But they are all very different from our proposed scheme. In detail, Lai’s definition of window-limited CUSUM statistics is $S_t = \max_{t-w < k < t} \sum_{i=k+1}^t \log \frac{f_0(\xi_i, \theta)}{f_\infty(\xi_i)}$, thus it also assumed full knowledge of the pre- and post-change distributions. Therefore, we must differentiate our algorithm from those in the literature that usually refers to the above statistics and may coincides with window-limited GLR in some senses, and is non-recursive, while our proposed WL-CUSUM enjoys recursive update and the estimator $\hat{\theta}$ is also differently constructed.

A. Assumptions and Useful Observations

Regarding the estimates $\{\hat{\theta}_t\}$ since it is not our intention to promote any specific estimation method, we will impose general characteristics that are enjoyed by most reasonable estimators (such as MLE). In particular, we make the following key assumptions for the estimator.

A1: Under the post-change regime P_0^θ , we assume that $E_0^\theta[\hat{\theta}_t] = \theta$ (unbiased estimator¹). If we write $\hat{\theta}_t = \theta + \mathcal{E}_t$, the zero mean estimation error has a covariance matrix of the form $E_0^\theta[\mathcal{E}_t \mathcal{E}_t^\top] = \frac{1}{w} \Sigma_0 (1 + o(1))$. Matrix Σ_0 is

of the order of a constant when considered as a function of w .

A2: Under the pre-change regime P_∞ , we assume that $E_\infty[\hat{\theta}_t] = \theta_\infty$. If we write $\hat{\theta}_t = \theta_\infty + \mathcal{E}_t$, the zero mean estimation error has a covariance matrix of the form $E_\infty[\mathcal{E}_t \mathcal{E}_t^\top] = \frac{1}{w} \Sigma_\infty (1 + o(1))$. Matrix Σ_∞ is of the order of a constant as a function of w .

With A1, we require our estimator, when applied to post-change data to provide reliable estimates of the correct parameter vector θ . We expect the quality of our estimate to improve with increasing window size w , since the error covariance matrix is inversely proportional to w . When applied to pre-change data, the estimator behavior is described by A2. We assume that it provides estimates close to some value θ_∞ , where θ_∞ is estimator dependent. For example, in the case of the MLE, we have that

$$\theta_\infty = \arg \max_{\theta} E_\infty[\log f_0(\xi_1, \theta)],$$

which is the limiting form of (7) after we normalize with the window size w and invoke the Law-of-Large Numbers. We note that the MLE satisfies A1 and A2 since MLE is asymptotically unbiased (satisfying the footnote 1) under mild conditions [50]. Let us now continue with our assumptions. The next assumption refers to the Kullback-Leibler (KL) information numbers and the second moment of the log-likelihood ratio.

A3: Consider the two KL information numbers and the second moment of the log-likelihood ratio under the P_0^θ measure

$$\begin{aligned} I_0 &= E_0^\theta \left[\log \frac{f_0(\xi_1, \theta)}{f_\infty(\xi_1)} \right], \\ I_\infty &= -E_\infty \left[\log \frac{f_0(\xi_1, \theta_\infty)}{f_\infty(\xi_1)} \right], \\ J_0 &= E_0^\theta \left[\left(\log \frac{f_0(\xi_1, \theta)}{f_\infty(\xi_1)} \right)^2 \right]. \end{aligned}$$

We assume that all three quantities are strictly *positive* and *bounded* for every $\theta \in \Theta$ of interest.

We see that I_∞ involves the parameter value θ_∞ , which is estimated when the data are under the pre-change regime. This information number can be strictly positive if, for example, the pre-change density $f_\infty(\cdot)$ cannot be expressed as the post-change density $f_0(\cdot, \theta_\infty)$ for some particular value $\theta_\infty \in \Theta$ that belongs to the allowable set of post-change parameter values. Unfortunately, the analysis that follows does not cover the case where $f_\infty(\cdot) = f_0(\cdot, \theta_\infty)$ and it requires significant technical modifications to address this particular possibility. On the other hand we can avoid the occurrence of this case by defining a suitable set Θ that does not contain θ_∞ .

In our derivations, we will encounter quantities that resemble the KL information numbers and the second moment of the log-likelihood ratio introduced in Assumption A3, but with the parameter θ replaced by an estimate. In particular, we will be interested in the following alternatives:

$$\hat{I}_0 = E_0^\theta \left[\log \frac{f_0(\xi_t, \hat{\theta}_{t-1})}{f_\infty(\xi_t)} \right],$$

¹In fact our analysis can also accommodate asymptotically unbiased estimators provided that the norm square of the bias is $o(1/w)$. It is for simplicity that we limit ourselves to the unbiased case.

$$\hat{l}_\infty = -E_\infty \left[\log \frac{f_0(\xi_t, \hat{\theta}_{t-1})}{f_\infty(\xi_t)} \right],$$

$$\hat{J}_0 = E_0^\theta \left[\left(\log \frac{f_0(\xi_t, \hat{\theta}_{t-1})}{f_\infty(\xi_t)} \right)^2 \right].$$

Due to Assumptions A1 and A2 we expect \hat{l}_0 , \hat{l}_∞ , and \hat{J}_0 to be close to their exact counterparts l_0 , l_∞ , J_0 . In fact we can specify their relationship more precisely by applying a Taylor expansion on $\log(f_0(\xi_t, \hat{\theta}_{t-1})/f_\infty(\xi_t))$ around the mean of $\hat{\theta}_{t-1}$ and retaining the first three terms. Additionally, we could make suitable assumptions on the smoothness of $f_0(\xi, \theta)$ to guarantee the effectiveness of these approximations. In order to avoid these common technicalities, we propose to simply assume that such an expansion is valid without more details. This will allow us to focus on the more interesting question of the WLCUSUM asymptotic optimality, which will require a number of novel results due to the w -dependency of the approximate log-likelihood ratios $\{\log(f_0(\xi_t, \hat{\theta}_{t-1})/f_\infty(\xi_t))\}$. Consequently, our last assumption expresses how \hat{l}_0 , \hat{l}_∞ , \hat{J}_0 are related to l_0 , l_∞ , J_0 .

A4 (Taylor Expansion Based Approximations): The quantities \hat{l}_0 , \hat{l}_∞ , and \hat{J}_0 can be written as

$$\begin{aligned}\hat{l}_0 &= l_0 - \frac{1}{2w} \text{trace}\{\Sigma_0 F_0\} (1 + o(1)), \\ \hat{l}_\infty &= l_\infty + \frac{1}{2w} \text{trace}\{\Sigma_\infty F_\infty\} (1 + o(1)), \\ \hat{J}_0 &= J_0 + \frac{1}{w} \text{trace}\{\Sigma_0 Q_0\} (1 + o(1)),\end{aligned}\quad (10)$$

where F_0, F_∞, Q_0 are matrices of the order of a constant with respect to the window size w .

By applying Taylor expansion on the approximate log-likelihood function, it is possible to identify the exact form of F_0, F_∞, Q_0 in (10) for unbiased estimators satisfying Assumptions A1 and A2. We present this in the following lemma.

Lemma 2: The matrices entering in (10) in Assumption A4 have the following form

$$\begin{aligned}F_0 &= E_0^\theta \left[\left(\frac{\nabla_\theta f_0(\xi_1, \theta)}{f_0(\xi_1, \theta)} \right) \left(\frac{\nabla_\theta f_0(\xi_1, \theta)}{f_0(\xi_1, \theta)} \right)^\top \right], \\ F_\infty &= E_\infty \left[\left(\frac{\nabla_\theta f_0(\xi_1, \theta_\infty)}{f_0(\xi_1, \theta_\infty)} \right) \left(\frac{\nabla_\theta f_0(\xi_1, \theta_\infty)}{f_0(\xi_1, \theta_\infty)} \right)^\top \right. \\ &\quad \left. - \frac{\nabla_{\theta\theta} f_0(\xi_1, \theta_\infty)}{f_0(\xi_1, \theta_\infty)} \right], \\ Q_0 &= F_0 + E_0^\theta \left[\left(\log \frac{f_0(\xi_1, \theta)}{f_\infty(\xi_1)} \right) \left\{ \frac{\nabla_{\theta\theta} f_0(\xi_1, \theta)}{f_0(\xi_1, \theta)} \right. \right. \\ &\quad \left. \left. - \left(\frac{\nabla_\theta f_0(\xi_1, \theta)}{f_0(\xi_1, \theta)} \right) \left(\frac{\nabla_\theta f_0(\xi_1, \theta)}{f_0(\xi_1, \theta)} \right)^\top \right\} \right],\end{aligned}$$

where $\nabla_{\theta\theta} f_0(\xi, \theta)$ denotes the Hessian of $f_0(\xi, \theta)$ with respect to θ .

Proof: Demonstrating the validity of these formulas presents no particular difficulty. We apply Taylor expansion on $\log(f_0(\xi_t, \hat{\theta}_{t-1})/f_\infty(\xi_t))$ with respect to $\hat{\theta}_{t-1}$ around its mean θ and retain the first three terms. Then we make use of the independence between ξ_t and $\hat{\theta}_{t-1}$ and therefore between ξ_t and \mathcal{E}_{t-1} (the estimation error of $\hat{\theta}_{t-1}$). The first order

term in the expansion has average 0 because \mathcal{E}_{t-1} is zero mean. Consequently, we end up with the expectation of the second order term with respect to \mathcal{E}_{t-1} while, as we assumed, higher order terms are regarded as negligible and captured by the $o(1)$ components in (10). In order to obtain the desired results we must also note that $E_0^\theta[\nabla_\theta f_0(\xi_1, \theta)/f_0(\xi_1, \theta)] = 0$, $E_0^\theta[\nabla_{\theta\theta} f_0(\xi_1, \theta)/f_0(\xi_1, \theta)] = 0$ and that for a deterministic matrix Q , we can write $E[\mathcal{E}_{t-1}^\top Q \mathcal{E}_{t-1}] = E[\text{trace}\{Q \mathcal{E}_{t-1} \mathcal{E}_{t-1}^\top\}] = \text{trace}\{Q E[\mathcal{E}_{t-1} \mathcal{E}_{t-1}^\top]\}$. Computations are straightforward, thus we omit further details. \square

Remark 4 (Examples of Gaussian Distribution): We provide a concrete example of the quantities above under Gaussian distributions. Assume one-dimensional Gaussian mean shift from $\mathcal{N}(0, 1)$ to $\mathcal{N}(\theta, 1)$ with post-change mean equal to 1, and assume the set $\Theta = \{\theta : \theta \geq 0.5\}$ as the possible post-change mean values. When using the maximum likelihood estimator, we have $\hat{\theta}_t = \frac{1}{w} \sum_{i=0}^{w-1} \xi_{t-i}$ if $\frac{1}{w} \sum_{i=0}^{w-1} \xi_{t-i} \geq 0.5$ and $\hat{\theta}_t = 0.5$ otherwise. Under the post-change regime, the asymptotic distribution of $\hat{\theta}_t$ is $\mathcal{N}(\theta, 1/w)$, while under the pre-change regime, $\hat{\theta}_t$ converges in probability to 0.5. In this case we have $\theta = 1$, $\theta_\infty = 0.5$, $l_0 = \theta^2/2$, $l_\infty = 0.5^2/2$, $J_0 = \theta^4/4 + \theta^2$, and Fisher information $F_0 = 1$, $F_\infty = 1$, $Q_0 = F_0 - \theta^2/2 = 1/2$, and $\hat{l}_0 \approx l_0 - 1/(2w)$, $\hat{l}_\infty = 1/8 + o(1/w)$, $\hat{J}_0 \approx J_0 + 1/(2w)$. It is worthwhile commenting that when the pre-change distribution is non-normal, the parameter set Θ will equal to \mathbb{R} and the MLE $\hat{\theta}_t$ is asymptotically normal under mild conditions for both the pre- and post-change regimes.

Let us now make some useful observations regarding the quantities we introduced above. From Equation (10) we deduce that the window size w must be larger than $\text{trace}\{\Sigma_0 F_0\}/l_0$ so that $\hat{l}_0 > 0$, a property which is necessary for successful detection. Indeed, only when $\hat{l}_0 > 0$, the WLCUSUM statistic will exhibit a positive drift after change, forcing the corresponding statistic to increase and finally exceed the positive threshold. Meanwhile, we require the pre-change drift $-\hat{l}_\infty$ to be negative; this holds in general for any w and any estimate as we discuss in detail in Section V.

We note from its definition that F_0 is the Fisher Information matrix. Since for any unbiased estimator we have validity of the Cramer-Rao Lower Bound, namely $\Sigma_0 \succcurlyeq (F_0)^{-1}$, we conclude that

$$\text{trace}\{\Sigma_0 F_0\} = \text{trace}\{F_0^{\frac{1}{2}} \Sigma_0 F_0^{\frac{1}{2}}\} \geq \text{trace}\{F_0^{\frac{1}{2}} F_0^{-1} F_0^{\frac{1}{2}}\} = K; \quad (11)$$

where, we recall that K is the dimension of the parameter vector θ . Therefore, the maximum likelihood estimate asymptotically maximizes the approximate KL information number \hat{l}_0 , thus resulting in the smallest detection delay among all consistent estimators.

Consider now the *mismatched* version of the KL information number $E_0^\theta[\log(f_0(\xi_1, \vartheta)/f_\infty(\xi_1))]$ with $\vartheta \neq \theta$, then

$$\begin{aligned}E_0^\theta \left[\log \frac{f_0(\xi_1, \vartheta)}{f_\infty(\xi_1)} \right] &= E_0^\theta \left[\log \frac{f_0(\xi_1, \theta)}{f_\infty(\xi_1)} \right] + E_0^\theta \left[\log \frac{f_0(\xi_1, \vartheta)}{f_0(\xi_1, \theta)} \right] \\ &\leq E_0^\theta \left[\log \frac{f_0(\xi_1, \theta)}{f_\infty(\xi_1)} \right] = l_0.\end{aligned}$$

This observation and the fact that l_0 is constant allows us to deduce that

$$\begin{aligned} \mathbb{E}_0^\theta \left[\log \frac{f_0(\xi_t, \hat{\theta}_{t-1})}{f_\infty(\xi_t)} \middle| \mathcal{F}_{t-1} \right] &\leq l_0 \\ \Rightarrow \hat{l}_0 = \mathbb{E}_0^\theta \left[\log \frac{f_0(\xi_t, \hat{\theta}_{t-1})}{f_\infty(\xi_t)} \right] &\leq l_0, \end{aligned} \quad (12)$$

which will be used in the derivations that follow.

V. THEORETICAL ANALYSIS

We are now ready to analyze the proposed WLCUSUM test. We begin by considering next the average false alarm.

Lemma 3 (ARL of WLCUSUM): The WLCUSUM defined in (6),(8) satisfies

$$\mathbb{E}_\infty[\mathcal{T}] \geq e^\nu.$$

Additionally, for the drift under the pre-change regime, we have $-\hat{l}_\infty = \mathbb{E}_\infty[\log \frac{f_0(\xi_t, \hat{\theta}_{t-1})}{f_\infty(\xi_t)}] < 0$.

Proof: We compute the average false alarm period using similar ideas as in [3, Lemma 8.2.1]. For $t > w$ let us define a Shiryaev-Roberts like statistic $\{\mathcal{L}_t\}$ through the recursion

$$\mathcal{L}_t = (\mathcal{L}_{t-1} + 1) \frac{f_0(\xi_t, \hat{\theta}_{t-1})}{f_\infty(\xi_t)}, \quad \mathcal{L}_w = 0.$$

Interestingly, $\{\mathcal{L}_t\}$ preserves the characteristic martingale property with respect to the P_∞ measure enjoyed by the classical Shiryaev-Roberts (SR) test statistic

$$L_t = (L_{t-1} + 1) \frac{f_0(\xi_t, \theta)}{f_\infty(\xi_t)}, \quad L_0 = 0,$$

even when we replace θ with the estimates $\{\hat{\theta}_t\}$. Indeed we observe that

$$\begin{aligned} \mathbb{E}_\infty[\mathcal{L}_t - t | \mathcal{F}_{t-1}] &= \mathbb{E}_\infty \left[(\mathcal{L}_{t-1} + 1) \frac{f_0(\xi_t, \hat{\theta}_{t-1})}{f_\infty(\xi_t)} - t \middle| \mathcal{F}_{t-1} \right] \\ &= \mathcal{L}_{t-1} - (t-1). \end{aligned}$$

The last equality is true because given \mathcal{F}_{t-1} we have $\hat{\theta}_{t-1}$ fixed, therefore $f_0(\xi_t, \hat{\theta}_{t-1})$ is a legitimate probability density for ξ_t . The martingale property of $\{\mathcal{L}_t - t\}$ and usage of Optional Sampling allows us to write for any stopping time T with finite expectation that

$$\mathbb{E}_\infty[\mathcal{L}_T - T] = \mathbb{E}_\infty[\mathcal{L}_w - w] = -w \Rightarrow \mathbb{E}_\infty[T] - w = \mathbb{E}_\infty[\mathcal{L}_T]. \quad (13)$$

Let us now recall the WLCUSUM update in (6) which after exponentiation can be equivalently written as

$$e^{S_t} = \max\{e^{S_{t-1}}, 1\} \frac{f_0(\xi_t, \hat{\theta}_{t-1})}{f_\infty(\xi_t)}, \quad e^{S_w} = 1.$$

Using induction, the fact that $e^{S_{w+1}} = \mathcal{L}_{w+1}$ and that for $x \geq 0$, $x+1 \geq \max\{x, 1\}$, it is straightforward to prove that for $t > w$ we have $\mathcal{L}_t \geq e^{S_t}$. This suggests that the SR-statistic is larger than the exponential of the WLCUSUM statistic. With this observation and using (13) we can now write

$$e^\nu \leq \mathbb{E}_\infty[e^{S_{\mathcal{T}}}] \leq \mathbb{E}_\infty[\mathcal{L}_{\mathcal{T}}] = \mathbb{E}_\infty[\mathcal{T}] - w \leq \mathbb{E}_\infty[\mathcal{T}],$$

which proves the desired inequality. We also observe that for any θ we have $\mathbb{E}_\infty[\log \frac{f_0(\xi_1, \theta)}{f_\infty(\xi_1)}] < 0$, therefore $\mathbb{E}_\infty[\log \frac{f_0(\xi_t, \hat{\theta}_{t-1})}{f_\infty(\xi_t)} | \mathcal{F}_{t-1}] < 0$, which in turn implies $\mathbb{E}_\infty[\log \frac{f_0(\xi_t, \hat{\theta}_{t-1})}{f_\infty(\xi_t)}] = -\hat{l}_\infty < 0$. \square

Equating the lower bound e^ν provided by Lemma 3 to the desired average false alarm period γ , assures that the false alarm constraint $\mathbb{E}_\infty[\mathcal{T}] \geq \gamma$ is satisfied. Consequently, the threshold we select to use is equal to

$$\nu = \log \gamma. \quad (14)$$

Consider now the negative drift $-\hat{l}_\infty$ mentioned in the lemma that appears under the pre-change regime. The drift under the P_∞ measure must be negative, because this assures restarts of the process and also that the average false alarm period will be an exponential function of the threshold. Since the drift is equal to $-\hat{l}_\infty$, we have that \hat{l}_∞ must be positive. As we argued in the proof of Lemma 3 the approximate KL information number \hat{l}_∞ is indeed positive, and using (10) we can study how the estimator affects this positive value. From (10) we can see that the first term of \hat{l}_∞ on the right-hand side is positive. One may wonder whether the second term due to the estimation error may also contribute to the positivity of \hat{l}_∞ . Of course, the sign of this term is estimator-dependent. However, in the case of the MLE, it is easy to see that $-F_\infty$ is the Hessian of $\mathbb{E}_\infty[\log f_0(\xi_1, \theta)]$ evaluated at the point θ_∞ where this function is maximized.² Consequently, $-F_\infty$ is negative definite and therefore F_∞ positive definite. This means that the sign of $\text{trace}\{F_\infty \Sigma_\infty\}$ will be positive as well, contributing to the positivity of \hat{l}_∞ and therefore the negativity of the drift. We must, however, emphasize that other estimators do not necessarily share this desirable property and positivity is assured for sufficiently large window w .

The next step in our analysis consists in computing the worst-case average detection delay of WLCUSUM. In the following lemma, we present an important property for our detection strategy, which is also shared by the classical CUSUM and considerably facilitates the performance computation.

Lemma 4 (Worst-Case Average Detection Delay): For any changepoint $\tau \geq 0$ we have that

$$\begin{aligned} \text{ess sup } \mathbb{E}_\tau^\theta[\mathcal{T} - \tau | \mathcal{T} > \tau, \mathcal{F}_\tau] \\ \leq w + \mathbb{E}_\tau^\theta \left[\mathbb{E}_\tau^\theta[(\mathcal{T}_\tau - \tau - w) \mathbb{1}_{\{\mathcal{T}_\tau > \tau + w\}} | \xi_{\tau+1}^{\tau+w}] \right] = \mathbb{E}_0^\theta[\mathcal{T}]. \end{aligned}$$

Proof: Consider first a change at $\tau = 0$. Since we start testing at $w+1$ we can write

$$\mathbb{E}_0^\theta[\mathcal{T}] = w + \mathbb{E}_0^\theta \left[\mathbb{E}_0^\theta[(\mathcal{T} - w) \mathbb{1}_{\{\mathcal{T} > w\}} | \xi_1^w] \right].$$

Suppose now that the change happens at τ then for $t > \tau + w$ it is clear that the test statistic S_t is larger than the test statistic $S_{t,\tau}$ generated by starting the WLCUSUM at time $\tau + w + 1$ and using the samples $\{\xi_{\tau+w}, \dots, \xi_{\tau+1}\}$ to form the first estimate $\hat{\theta}_{\tau+w}$. This suggests that if we use $S_{t,\tau}$ instead of S_t in (8) then we will stop at a time \mathcal{T}_τ that satisfies $\mathcal{T}_\tau \geq \mathcal{T}$. Clearly we also have that \mathcal{T}_τ is independent from \mathcal{F}_τ , because

²Provided of course that this value is not on the border of the allowable parameter set Θ and corresponds to an unconstrained optimizer.

in $\mathcal{S}_{t,\tau}$ the data \mathcal{F}_τ are not being used. With these observations in mind we can write

$$\begin{aligned} E_\tau^\theta[\mathcal{T} - \tau | \mathcal{T} > \tau, \mathcal{F}_\tau] &\leq w + E_\tau^\theta[(\mathcal{T} - \tau - w) \mathbb{1}_{\{\mathcal{T} > \tau + w\}} | \mathcal{T} > \tau, \mathcal{F}_\tau] \\ &= w + E_\tau^\theta \left[E_\tau^\theta[(\mathcal{T} - \tau - w) \mathbb{1}_{\{\mathcal{T} > \tau + w\}} | \xi_{\tau+1}^{\tau+w}, \mathcal{F}_\tau] | \mathcal{T} > \tau, \mathcal{F}_\tau \right] \\ &\leq w + E_\tau^\theta \left[E_\tau^\theta[(\mathcal{T}_\tau - \tau - w) \mathbb{1}_{\{\mathcal{T}_\tau > \tau + w\}} | \xi_{\tau+1}^{\tau+w}, \mathcal{F}_\tau] | \mathcal{T} > \tau, \mathcal{F}_\tau \right] \\ &= w + E_\tau^\theta \left[E_\tau^\theta[(\mathcal{T}_\tau - \tau - w) \mathbb{1}_{\{\mathcal{T}_\tau > \tau + w\}} | \xi_{\tau+1}^{\tau+w}] \right] = E_0^\theta[\mathcal{T}]. \end{aligned}$$

The last equality is true because starting the procedure at $\tau + w + 1$ when the change occurs at τ does not employ any information from \mathcal{F}_τ , consequently it is independent from \mathcal{F}_τ . Therefore, statistically this is the same as starting at $w + 1$ with the change occurring at 0. \square

The analysis of $E_0^\theta[\mathcal{T}]$ is not simple due to the particular updating rule of \mathcal{S}_t . Since we are only interested in finding a suitable upper bound, we are going to introduce an alternative stopping time \mathcal{T}' , which is easier to analyze. Its definition and connection to \mathcal{T} are presented in the following lemma.

Lemma 5: Assume that the observations follow the post-change regime and that $\hat{l}_0 > 0$. For $t > w$ define the process $\{\mathcal{U}_t\}$ using the recursion

$$\mathcal{U}_t = \mathcal{U}_{t-1} + \log \frac{f_0(\xi_t, \hat{\theta}_{t-1})}{f_\infty(\xi_t)}, \quad \mathcal{U}_w = 0, \quad (15)$$

and the stopping time

$$\mathcal{T}' = \inf\{t > w : \mathcal{U}_t \geq \nu\},$$

then \mathcal{T}' stops a.s. and we have $\mathcal{T}' \geq \mathcal{T}$.

Proof: Because $\mathcal{S}_{w+1} = \mathcal{U}_{w+1} = \log \frac{f_0(\xi_{w+1}, \hat{\theta}_w)}{f_\infty(\xi_{w+1})}$ and $x^+ = \max\{x, 0\} \geq x$, using induction it is simple to prove that $\mathcal{S}_t \geq \mathcal{U}_t$ for $t > w$. This of course implies that \mathcal{U}_t will require more time than \mathcal{S}_t to reach the same threshold ν , which means that $\mathcal{T}' \geq \mathcal{T}$. Consequently, $E_0^\theta[\mathcal{T}'] \geq E_0^\theta[\mathcal{T}]$. The fact that \mathcal{T}' will stop a.s. is guaranteed by the positivity of $\hat{l}_0 = E_0^\theta[\log \frac{f_0(\xi_t, \hat{\theta}_{t-1})}{f_\infty(\xi_t)}]$. \square

Finding a bound for $E_0^\theta[\mathcal{T}']$ is simpler because as we see from its definition in (15) we have that \mathcal{U}_t is a sum of stationary (but w -dependent) terms. The estimate we are looking for is provided in the next theorem which identifies, with the help of Lemma 5, an upper bound for the performance of the WLCUSUM test.

Theorem 1 (WADD of WLCUSUM): Assume $\hat{l}_0 > 0$. We have the following upper bound for the worst-case performance of WLCUSUM

$$\begin{aligned} E_0^\theta[\mathcal{T}] &\leq E_0^\theta[\mathcal{T}'] \\ &\leq \frac{\log \gamma + \frac{\hat{l}_0}{l_0} + \left(\frac{\hat{l}_0}{l_0} \log \gamma\right)^{\frac{1}{2}} + w l_0 + \left(\frac{\hat{l}_0}{l_0} l_0 w\right)^{\frac{1}{2}}}{\hat{l}_0}, \quad (16) \end{aligned}$$

where \hat{l}_0 and \hat{J}_0 were defined in Assumption A4.

Proof: The complete proof of our claim involves several steps. We provide a proof sketch in three steps. First, from Lemma 4 we have that the worst-case average detection delay equals to $E_0^\theta[\mathcal{T}]$, thus we only need to bound $E_0^\theta[\mathcal{T}']$. Second,

we apply the Wald's identity for w -dependent samples to bound the expected stopping time $E_0^\theta[\mathcal{T}']$ for process $\{\mathcal{U}_t\}$ in (15), this is an upper bound for the detection delay $E_0^\theta[\mathcal{T}]$ based on Lemma 5. Finally, we substitute the threshold $\nu = \log \gamma$ according to Lemma 3 to ensure the ARL satisfies $E_\infty[\mathcal{T}] \geq \gamma$. Details can be found in the Appendix. \square

Remark 5: If we fix the window size w then from Theorem 1 we have that the upper bound of the worst-case average detection delay of WLCUSUM can be written as

$$E_0^\theta[\mathcal{T}] \leq \frac{\log \gamma}{\hat{l}_0} (1 + o(1)).$$

Because from (12) we have $\hat{l}_0 < l_0$ this means that we cannot assure first-order asymptotic optimality with fixed w . From (10) we see that the difference between the optimum denominator l_0 (enjoyed by the exact CUSUM) and \hat{l}_0 obtained by our proposed stopping time is (asymptotically) equal to $(1/w) \text{trace}\{\Sigma_0 F_0\}$. As we argued in (11) the factor $\text{trace}\{\Sigma_0 F_0\}$ due to the Cramer-Rao Lower Bound cannot become smaller than the length K of the parameter vector, which corresponds to the best possible performance since it maximizes the denominator \hat{l}_0 . We recall that this optimal value is attained asymptotically by the MLE.

In order for WLCUSUM to be first-order asymptotically optimum, it is necessary to increase w with the average run length constraint γ in a manner that guarantees that w remains negligible compared to the optimum CUSUM performance but grows sufficiently fast so that it provides efficient estimates of θ . Since the CUSUM detection delay increases as $\log \gamma$ (see Lemma 1) this means that we must select $w = o(\log \gamma)$. At the same time, to obtain improved estimates, we must let $w \rightarrow \infty$ as $\gamma \rightarrow \infty$. These observations will be taken into account to decide the appropriate growth rate of w .

From Lemma 1 we deduce that the classical CUSUM stopping time \mathcal{T} satisfies

$$E_0^\theta[\mathcal{T}] / \frac{\log \gamma}{l_0} = 1 + \Theta\left(\frac{1}{\log \gamma}\right).$$

Since CUSUM is (strictly) optimum [13], we have

$$E_0^\theta[\mathcal{T}] / \frac{\log \gamma}{l_0} \leq E_0^\theta[\mathcal{T}'] / \frac{\log \gamma}{l_0}.$$

Using the results of Theorem 1 and properly selecting the increase rate of w , in the next theorem, we demonstrate that WLCUSUM enjoys the desired first-order asymptotic optimality.

Theorem 2 (Asymptotic Optimality of WLCUSUM): If the window size w satisfies $w \rightarrow \infty$ as $\gamma \rightarrow \infty$ with $w = o(\log \gamma)$, then

$$E_0^\theta[\mathcal{T}] / \frac{\log \gamma}{l_0} \leq 1 + \Theta\left(\frac{1}{\sqrt{\log \gamma}}\right) + \Theta\left(\frac{1}{w}\right) + \Theta\left(\frac{w}{\log \gamma}\right).$$

The maximal convergence speed to unity is $1 + \Theta\left(\frac{1}{\sqrt{\log \gamma}}\right)$ and is achieved when $w = \Theta(\sqrt{\log \gamma})$.

Proof: Using (10) we note that the ratio \hat{J}_0/\hat{l}_0 in the numerator in (16) is a function of w , which for sufficiently large w can be bounded by a constant that does not depend on w . If we select $w = o(\log \gamma)$ but $w \rightarrow \infty$ as $\gamma \rightarrow \infty$,

replace it in the upper bound in (16) and also use (10) for the denominator then we can write

$$\frac{E_0^\theta[\mathcal{T}]}{\frac{\log \gamma}{l_0}} \leq \frac{1 + \Theta\left(\frac{1}{\log \gamma}\right) + \Theta\left(\frac{1}{\sqrt{\log \gamma}}\right) + \Theta\left(\frac{w}{\log \gamma}\right) + \Theta\left(\frac{w^{1/2}}{\log \gamma}\right)}{1 + \Theta\left(\frac{1}{w}\right)}. \quad (17)$$

All $\Theta(\cdot)$ terms in the right-hand side converge to 0 and the overall convergence rate is dominated by the slowest term among: $\{\Theta(\frac{1}{w}), \Theta(\frac{1}{\sqrt{\log \gamma}}), \Theta(\frac{w}{\log \gamma})\}$. Note that the best possible convergence rate towards 1 is obtained when we select $w = \Theta(\sqrt{\log \gamma})$. Therefore $w \sim \sqrt{\log \gamma}$ is the optimal choice but any window size w satisfying $w \rightarrow \infty$ as $\gamma \rightarrow \infty$ with $w = o(\log \gamma)$ is sufficient to guarantee the first-order optimality. \square

The immediate consequence of Theorem 2 is the asymptotic optimality of WLCUSUM since

$$1 = \lim_{\gamma \rightarrow \infty} E_0^\theta[\mathcal{T}] / \frac{\log \gamma}{l_0} \leq \lim_{\gamma \rightarrow \infty} E_0^\theta[\mathcal{T}] / \frac{\log \gamma}{l_0} \leq \lim_{\gamma \rightarrow \infty} \left\{ 1 + \Theta\left(\frac{1}{\sqrt{\log \gamma}}\right) + \Theta\left(\frac{1}{w}\right) + \Theta\left(\frac{w}{\log \gamma}\right) \right\} = 1,$$

when selecting $w = o(\log \gamma)$ with $w \rightarrow \infty$ as $\gamma \rightarrow \infty$, which proves that $E_0^\theta[\mathcal{T}] / \frac{\log \gamma}{l_0}$ attains the same nonzero limit as the optimum $E_0^\theta[\mathcal{T}] / \frac{\log \gamma}{l_0}$ and is therefore asymptotically optimum as well. As we have indicated, the optimal convergence rate towards 1 is of the form $\Theta((\log \gamma)^{-\frac{1}{2}})$ using the optimum window size $w = \Theta(\sqrt{\log \gamma})$. Although there is a definite loss in performance as compared to the rate of the exact CUSUM which is $\Theta((\log \gamma)^{-1})$ (Lemma 1), our scheme still enjoys asymptotic optimality.

Remark 6: One may claim that the specific rate we obtained is a consequence of the crude upper bound $O(\sqrt{\nu})$ we established for the average overshoot (see Appendix, proof of Theorem 1), and we could have obtained better results if this bound were a constant as in the i.i.d. case [51]. In fact a constant would have resulted in a term of the form $\Theta(\frac{1}{\log \gamma})$ in place of the $\Theta(\frac{1}{\sqrt{\log \gamma}})$ we have now in the numerator of (17). For the computation of the convergence rate, this would have led to the smallest power satisfying $\beta = \min\{\alpha, 1 - \alpha\}$ instead of $\beta = \min\{\alpha, 1 - \alpha, \frac{1}{2}\}$. However, if we maximize $\min\{\alpha, 1 - \alpha\}$ we still obtain that the best power rate is equal to $\frac{1}{2}$ which is achieved when $\alpha = \frac{1}{2}$ or equivalently $w = \Theta(\sqrt{\log \gamma})$.

Remark 7 (Comparison With Window-Limited GLR): In order for the window-limited GLR procedure in (9) to enjoy first-order asymptotic optimality, the window size w must also grow with γ in a certain rate. In particular w must satisfy $\liminf w / \log \gamma > l_0^{-1}$ and $\log w = o(\log \gamma)$ to ensure effective detection [14], namely window at least as large as the CUSUM average detection delay. In the proposed WLCUSUM procedure, we only require a window size $w = o(\log \gamma)$ with $w \rightarrow \infty$ as $\gamma \rightarrow \infty$, which can be much smaller, even negligible, than the requirement in the window-limited GLR.

In addition to the optimal order $w = \Theta(\sqrt{\log \gamma})$ obtained in Theorem 2, we also provide a detailed characterization of the constant terms, which will be helpful for choosing the

appropriate window size in practice. It should be mentioned that when the post-change quantities are unknown in practice, we can use the parallel variant in Section VI instead.

Lemma 6 (Optimal Window Size): The optimal window size that minimizes the upper bound of the WADD of WLCUSUM given in (16) is

$$w = \frac{\sqrt{\text{trace}\{\Sigma_0 F_0\}}}{l_0 \sqrt{2}} \sqrt{\log \gamma} \left\{ 1 + \Theta\left(\frac{1}{(\log \gamma)^{\frac{1}{4}}}\right) \right\}.$$

Proof: Substitute (10) for \hat{l}_0 and \hat{J}_0 in (16) and note that the upper bound in (16) can be written as

$$\frac{\log \gamma}{l_0} \left\{ 1 + \left(\frac{J_0}{l_0}\right)^{\frac{1}{2}} \frac{1}{(\log \gamma)^{\frac{1}{2}}} + \frac{l_0}{\log \gamma} w + \frac{\text{trace}\{\Sigma_0 F_0\}}{l_0} \frac{1}{2w} + \Theta\left(\frac{w^{\frac{1}{2}}}{\log \gamma}\right) + \Theta\left(\frac{1}{w^2}\right) \right\}.$$

Given from Theorem 2 that the optimum w we are seeking satisfies $w = \Theta(\sqrt{\log \gamma})$, the previous expression can be further simplified as follows

$$\begin{aligned} & \frac{\log \gamma}{l_0} \left\{ 1 + \left(\frac{J_0}{l_0}\right)^{\frac{1}{2}} \frac{1}{(\log \gamma)^{\frac{1}{2}}} + \frac{l_0}{\log \gamma} w + \frac{\text{trace}\{\Sigma_0 F_0\}}{l_0} \frac{1}{2w} + w^{\frac{1}{2}} \Theta\left(\frac{1}{\log \gamma}\right) \right\} \\ & \leq \frac{\log \gamma}{l_0} \left\{ 1 + \left(\frac{J_0}{l_0}\right)^{\frac{1}{2}} \frac{1}{(\log \gamma)^{\frac{1}{2}}} + \frac{l_0}{\log \gamma} w + \frac{\text{trace}\{\Sigma_0 F_0\}}{l_0} \frac{1}{2w} + \frac{w^{\frac{1}{2}} D}{\log \gamma} \right\}, \end{aligned}$$

for some constant D coming from the definition of $\Theta(\cdot)$. Minimizing the last expression over w yields the following equation for the derivative

$$\frac{l_0}{\log \gamma} - \frac{\text{trace}\{\Sigma_0 F_0\}}{l_0} \frac{1}{2w^2} + \frac{1}{2w^{\frac{1}{2}} \log \gamma} D = 0.$$

Finding a more precise form for the optimum w since we know its order of magnitude, amounts to selecting $w = C\sqrt{\log \gamma}(1 + \Theta((\log \gamma)^{-\rho}))$, $\rho > 0$ and specifying C and ρ . Substituting into the previous equation we obtain

$$\begin{aligned} & \frac{l_0}{\log \gamma} - \frac{\text{trace}\{\Sigma_0 F_0\}}{2l_0 C^2 \log \gamma} \left(1 - \Theta\left(\frac{1}{(\log \gamma)^\rho}\right) \right) + \Theta\left(\frac{1}{(\log \gamma)^{\frac{5}{4}}}\right) \\ & = \frac{l_0}{\log \gamma} - \frac{\text{trace}\{\Sigma_0 F_0\}}{2l_0 C^2 \log \gamma} + \Theta\left(\frac{1}{(\log \gamma)^{\frac{5}{4}}}\right) - \Theta\left(\frac{1}{(\log \gamma)^{1+\rho}}\right) \\ & = 0, \end{aligned}$$

where for $|\epsilon| \ll 1$ we used the approximation $1/(1 + \epsilon)^2 \approx 1 - 2\epsilon$. As we can see there are two terms which are of different order of magnitude therefore we need to make them 0 separately. This means

$$\begin{aligned} & \frac{l_0}{\log \gamma} - \frac{\text{trace}\{\Sigma_0 F_0\}}{l_0} \frac{1}{2C^2 \log \gamma} = 0, \\ & \Theta\left(\frac{1}{(\log \gamma)^{\frac{5}{4}}}\right) - \Theta\left(\frac{1}{(\log \gamma)^{1+\rho}}\right) = 0. \end{aligned}$$

The first results in $C = \sqrt{\text{trace}\{\Sigma_0 F_0\}}/l_0 \sqrt{2}$ and the second in $\rho = \frac{1}{4}$. This completes the proof of the lemma. \square

VI. PARALLEL WLCUSUM

Although the proposed WLCUSUM procedure in (6) and (8) enjoys first-order optimality, there are some practical concerns: (i) for any fixed w , the WLCUSUM procedure will start after observing w samples, which leads to a delay at least w ; (ii) the optimal window size derived in Lemma 6 might be unknown when the post-change signal strength l_0 is unknown; of course we can impose a pre-set lower bound on l_0 and design the window size according to this lower bound, but this may lead to larger than necessary detection delays. In this section, we provide a variant to WLCUSUM that can resolve the above issue being also suitable for practical implementation.

In case the optimal window size is difficult to estimate beforehand, we propose to perform *parallel* WLCUSUMs with different window sizes. More specifically, we run in parallel multiple WLCUSUMs with window sizes ranging from 1 up to some maximal value W . For each window size w , we denote the corresponding test statistic with $S_t(w)$. At each time instant t , all statistics are compared to a common threshold ν and the first time \mathcal{T}_P any of the statistics hits or exceeds ν is the time the parallel WLCUSUM will stop. If to each statistic we associate the corresponding stopping time $\mathcal{T}(w)$, then it is also true that the parallel WLCUSUM satisfies

$$\mathcal{T}_P = \min_{1 \leq w \leq W} \{\mathcal{T}(w)\}. \quad (18)$$

In adopting this approach, we do not wait for W samples in order to start the detection procedure. Indeed the first WLCUSUM with $w = 1$ has to wait for only one sample and then perform tests continuously. Furthermore, we do not need to specify an exact window size beforehand. In a sense by running in parallel multiple windows, it is mostly the best window that tends to be the first to stop.

The maximal window W does not have to be different from what we estimated in our previous analysis. Also, the average detection delay at $\tau = 0$ is still the worst case compared with any other change time $\tau > 0$, i.e., the analysis in Lemma 4 can be modified to cover \mathcal{T}_P as well and we can easily demonstrate that the parallel WLCUSUM, with a suitable maximal window size W , is also first-order asymptotically optimal matching the performance of WLCUSUM with the optimal window. We provide the necessary proof of this claim in the next lemma and the discussion that follows.

Lemma 7: The average run length of the Parallel WLCUSUM with maximum window size W satisfies

$$E_\infty[\mathcal{T}_P] \geq \frac{1}{W} e^\nu.$$

Proof: Recall that $S_t(w)$ denotes the statistic of the WLCUSUM with window size w . Similarly to Lemma 3, let us define a Shiryaev-Roberts like statistic $\{\mathcal{L}_t(w)\}$, under a window size w , through the recursion

$$\mathcal{L}_t(w) = (\mathcal{L}_{t-1}(w) + 1) \frac{f_0(\xi_t, \hat{\theta}_{t-1})}{f_\infty(\xi_t)}, \quad \mathcal{L}_w(w) = 0, \quad t \geq w + 1.$$

Following the arguments in Lemma 3, we have $\mathcal{L}_t(w) \geq e^{S_t(w)}$ for $t > w$ and $\{\mathcal{L}_t(w)\}$ preserves the characteristic martingale property with respect to the P_∞ measure. Applying Optional Sampling, we have $E_\infty[\mathcal{L}_{\mathcal{T}(w)} - T] = -w$ for

any stopping time T with finite expectation. The parallel WLCUSUM statistic can be written as $\max_{1 \leq w \leq W} \{S_t(w)\}$ with the corresponding stopping time satisfying $\mathcal{T}_P = \inf\{t > 0 : \max_{1 \leq w \leq W} \{S_t(w)\} \geq \nu\}$. For the stopping time we have

$$e^\nu \leq e^{\max_{1 \leq w \leq W} \{S_{\mathcal{T}_P}(w)\}} \leq \max_{1 \leq w \leq W} \{\mathcal{L}_{\mathcal{T}_P}(w)\} \leq \sum_{w=1}^W \mathcal{L}_{\mathcal{T}_P}(w).$$

Taking expectations under P_∞ on both sides yields

$$e^\nu \leq \sum_{w=1}^W (E_\infty[\mathcal{T}_P] - w) < W E_\infty[\mathcal{T}_P],$$

which proves the desired inequality. \square

Equating the lower bound $\frac{1}{W} e^\nu$ provided by Lemma 7 to the desired average false alarm period $\gamma > 1$ assures that the false alarm constraint $E_\infty[\mathcal{T}_P] \geq \gamma$ is satisfied. Consequently, the threshold we select to use is equal to

$$\nu = \log(W\gamma) \quad (19)$$

as opposed to $\log \gamma$ we use for a single window.

For the average detection delay, note that the delay of the parallel WLCUSUM will certainly be no larger than the delay of each individual WLCUSUM procedure. Therefore, using (18) and Theorem 1 we can write

$$\begin{aligned} E_0^\theta[\mathcal{T}_P] &\leq \min_{1 \leq w \leq W} E_0^\theta[\mathcal{T}(w)] \\ &\leq \min_{1 \leq w \leq W} \frac{\log(W\gamma) + \frac{\hat{J}_0}{l_0} + \left(\frac{\hat{J}_0}{l_0} \log(W\gamma)\right)^{\frac{1}{2}} + w l_0 + \left(\frac{\hat{J}_0}{l_0} l_0 w\right)^{\frac{1}{2}}}{\hat{l}_0}, \end{aligned}$$

where \hat{J}_0, \hat{l}_0 are the estimated information numbers under window size w . Comparing the previous upper bound to the one obtained in (16) we can see that if $\log W = o(\sqrt{\log \gamma})$ the two upper bounds are asymptotically equivalent for each w . This implies that the parallel WLCUSUM performs the optimization with respect to w *automatically* without the need of prior manual selection. Of course this is true as long as W exceeds the value of the optimum window and at the same time conforms with the requirement $\log W = o(\sqrt{\log \gamma})$. Regarding the last constraint it is trivially satisfied by the window sizes we employ namely $W = \Theta((\log \gamma)^\alpha)$. Indeed if we consider $W = \Theta((\log \gamma)^\alpha)$ with $1 > \alpha > \frac{1}{2}$ then $\log W = \Theta(\log \log \gamma) = o(\sqrt{\log \gamma})$ but also W exceeds the optimum window size since the latter satisfies $\Theta(\sqrt{\log \gamma})$.

Remark 8: The computational complexities of WLCUSUM, parallel WLCUSUM, and window-limited GLR can be ordered as follows:

WLCUSUM < parallel WLCUSUM < window-limited GLR.

Although all three procedures with the appropriate choice of window size are first-order asymptotically optimal, they differ in the number of operations needed per update. To form a more precise idea about the requirements in computation, let us assume that the estimation process needs constant complexity per estimate when considering estimates of windows of consecutive sizes. This is because it can compute each estimate by updating the estimate of the previous size. This is for example the case when computing arithmetic averages

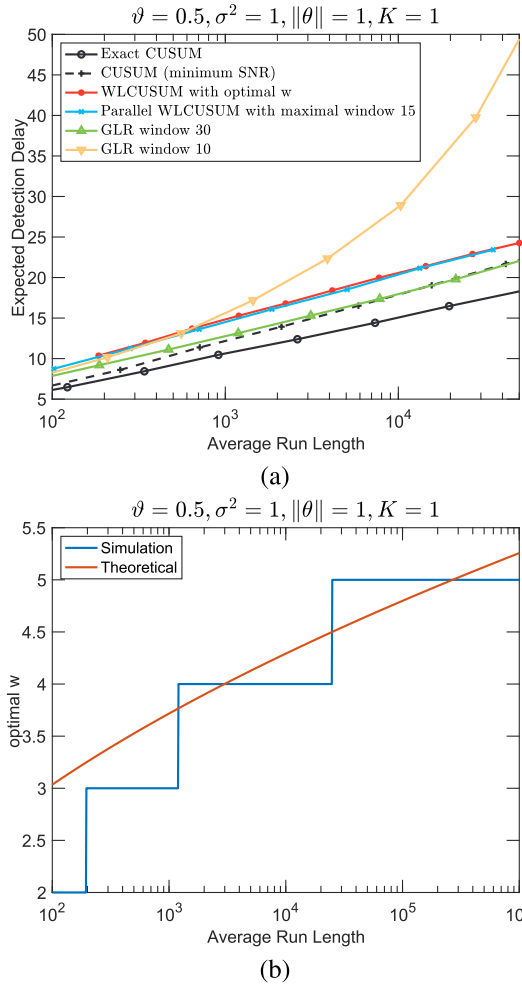


Fig. 1. (a) Detection delays for univariate normal mean-shift from 0 to $\theta = 1$, barrier $\vartheta = 0.5$; (b) Corresponding optimal window size as a function of ARL for WLCUSUM.

of samples of increasing windows since these estimates can be performed recursively. For the same estimation problem if we are interested only in a single window then the complexity of the estimator is proportional to the window size. Of course complexity can be significantly higher when estimates need to be obtained through iterative solutions.

Let us start with WLCUSUM with window size w . Its complexity is $\Theta(w)$ since we compute the estimate with complexity $\Theta(w)$ and then update the corresponding statistic with a constant number of operations. The parallel WLCUSUM with maximal window W will require $\Theta(W)$ operations to compute all the estimates of *consecutive* window sizes and $\Theta(W)$ operations to update the corresponding statistics (constant number of operations per statistic), therefore the total complexity is still $\Theta(W)$. In the case of the window-limited GLR if w is the size of the adopted window then we need $\Theta(w)$ to compute the estimates of the consecutive smaller window sizes. Then, each such estimate is applied to all the samples in the smaller window to form an approximation of the corresponding log-likelihood ratio. The total complexity required by these updates is $\Theta(1) + \dots + \Theta(w) = \Theta(w^2)$. The overall complexity for estimation and updates is $\Theta(w) + \Theta(w^2) = \Theta(w^2)$.

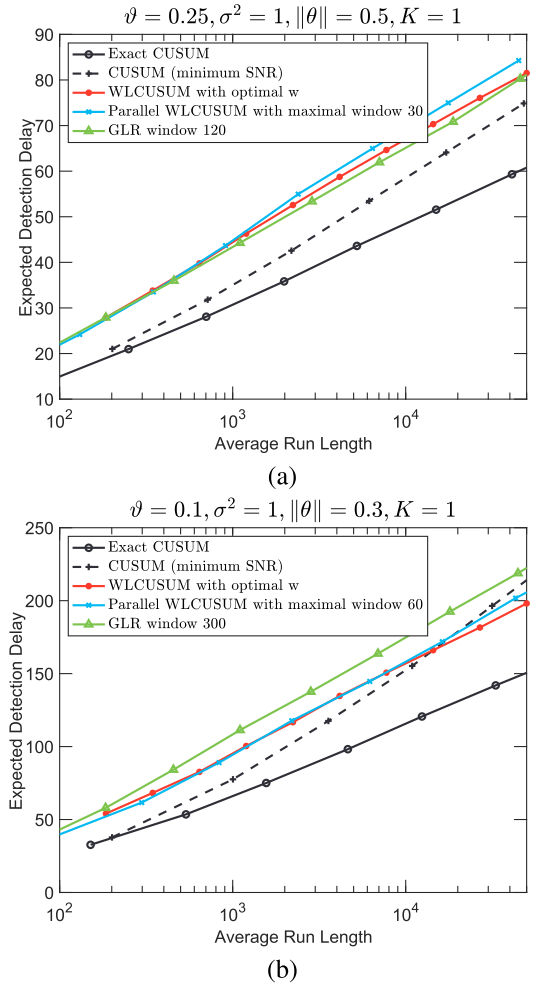


Fig. 2. Detection delays of different methods for univariate normal mean-shift with (a) $\theta = 0.5$ and (b) $\theta = 0.3$.

For WLCUSUM if we consider the optimal window size $w = \Theta(\sqrt{\log \gamma})$, this results in complexity $\Theta(\sqrt{\log \gamma})$ per time update. For the parallel WLCUSUM with a maximal window size $W = \Theta((\log \gamma)^\alpha)$ and $1 > \alpha > \frac{1}{2}$ to cover the optimal window size, we need complexity $\Theta((\log \gamma)^\alpha)$. In the case of the window-limited GLR we know that w must be at least of the same order as the detection delay therefore $w = \Theta(\log \gamma)$, suggesting that the corresponding computational complexity is $\Theta((\log \gamma)^2)$. As we can see the difference in number of operations between the two versions of WLCUSUM and the window-limited GLR is significant.

VII. EXPERIMENTS

Simulation studies are performed to compare WLCUSUM and parallel WLCUSUM to the exact CUSUM and the window-limited GLR. In the exact CUSUM all parameters are considered known and the corresponding detection delay is optimum [13]. In all experiments, the average detection delay and average false alarm period are obtained using direct estimation and averaging over 1000 independent trials.

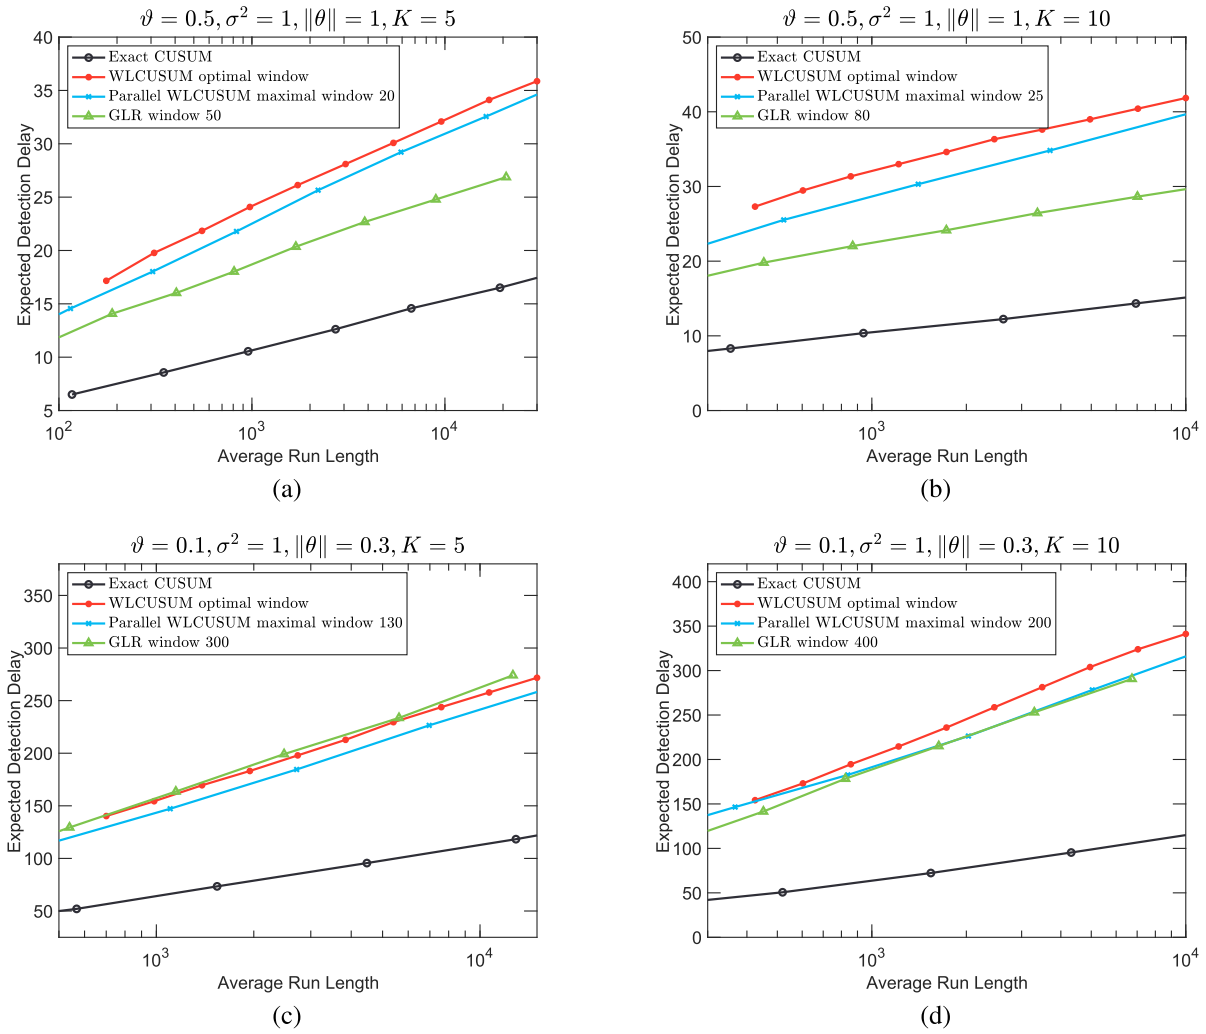


Fig. 3. Detection delay for multivariate normal mean-shift from 0 to θ with (a) $\|\theta\| = 1$, $\vartheta = 0.5$, $K = 5$; (b) $\|\theta\| = 1$, $\vartheta = 0.5$, $K = 10$; (c) $\|\theta\| = 0.3$, $\vartheta = 0.1$, $K = 5$; (d) $\|\theta\| = 0.3$, $\vartheta = 0.1$, $K = 10$.

A. Univariate Normal Mean-Shift

We consider a normal mean-shift example. Let $f_\infty(\cdot)$ be a canonical Gaussian pdf $\mathcal{N}(0, I_K)$ and $f_0(\cdot, \theta)$ a Gaussian $\mathcal{N}(\theta, I_K)$. Since $f_0(\cdot, \theta)$ with $\theta = 0$ is equal to $f_\infty(\cdot)$ we impose a constraint on θ in order to avoid the two densities becoming equal. For the mean θ we assume that $\theta \in \Theta = \{\theta : \|\theta\| \geq \vartheta > 0\}$ where the *barrier* value ϑ is considered known and can be interpreted as the smallest signal strength that we are interested in detecting. For our estimate of the mean we employ the MLE with the solution projected onto the allowable parameter set Θ . Specifically

$$\hat{\theta}_t = \frac{\bar{\theta}_t}{\|\bar{\theta}_t\|} \max\{\|\bar{\theta}_t\|, \vartheta\}, \quad \text{where } \bar{\theta}_t = \frac{1}{w} \sum_{s=0}^{w-1} \xi_{t-s}.$$

We first present a suite of numerical results for the one-dimensional ($K = 1$) case with varying signal strengths. In the first setting, we set $\theta = 1$ and $\vartheta = 0.5$. Figure 1(a) depicts the worst-case expected detection delay versus the average false alarm period (ARL), where the black line corresponds to the exact CUSUM procedure which lies below all the other curves. The red line corresponds to WLCUSUM with

optimal window size with Figure 1(b) depicting how the optimal window must change as a function of ARL (theoretical and simulations). In Figure 1(a) with blue we can also see the parallel WLCUSUM procedure with maximal window size $W = 15$. For comparison we also present in green the window-limited GLR defined in (9) with a window size of 30 (chosen to be larger than $\log \text{ARL}/l_0$ as explained in Remark 7) and in yellow with an insufficient window size 10. In the latter we notice the remarkable performance degradation if we do not use the appropriate window. Regarding our proposed methods we observe that the parallel WLCUSUM has nearly similar performance as the WLCUSUM with optimal window size, without requiring any prior specification of the optimal window. The window-limited GLR (with sufficient window size), in this scenario exhibits a smaller detection delay than both versions of WLCUSUM but at the expense of a higher computational cost as explained in Remark 8. Finally, we compare all procedures with the CUSUM test under minimal signal strength in black dashed line, i.e., we calculate the CUSUM test using densities of $\mathcal{N}(0.5, 1)$ and $\mathcal{N}(0, 1)$.

For the optimal window size depicted in Figure 1(b) we simulated sizes varying from 1 to 15 and for each ARL value

we selected the size with the smallest detection delay. This is depicted by the blue curve. With red we have plotted the theoretical optimal value obtained in Lemma 6. We observe a relatively good agreement between the two possibilities for large ARL values which is to be expected since our theoretical results are asymptotic.

In Figure 2(a) and (b) we performed similar experiments but under more difficult detection scenarios. In particular in (a) we consider the change in the mean from 0 to $\theta = 0.5$ with a barrier value of $\vartheta = 0.25$ while in (b) from 0 to $\theta = 0.3$ with barrier $\vartheta = 0.1$. Again both versions of WLCUSUM have comparable performance however now the window-limited GLR, despite its high complexity as a result of the corresponding large windows, does not enjoy better performance than WLCUSUM. In fact in the second more challenging detection example it is clearly worse than WLCUSUM. In addition, we see that the CUSUM under minimum signal strength also becomes much worse as ARL becomes large and is worse than WLCUSUM in the second more challenging detection example.

B. Varying Dimensions

To illustrate the performance on multivariate distributions, we performed simulations with varying parameter dimensions. As shown in (11), the dimension of θ will affect the appropriate choice of window size as well as the quantities \hat{l}_0 , l_∞ , and \hat{j}_0 . We consider two dimensions: $K = 5$ and $K = 10$. In both cases we simulate two different mean changes (i) from 0 to θ with $\|\theta\| = 1$ and barrier $\vartheta = 0.5$, and (ii) from 0 to θ with $\|\theta\| = 0.3$ and $\vartheta = 0.1$. The appropriate range of window sizes is chosen based on Lemma 6.

Figure 3(a) and (b) depicts the detection delay of the competing schemes under a large change namely from 0 to θ with $\|\theta\| = 1$ and barrier value $\vartheta = 0.5$ for the two lengths $K = 5, 10$. We can see in Figure 3(a) and (b) that in both cases the window-limited GLR exhibits a smaller detection delay compared to WLCUSUM but, as we explained, at the expense of a significantly higher computational complexity. However this advantage tends to disappear when the detection problem becomes more challenging as in the case of the small change in the mean from 0 to $\|\theta\| = 0.3$ with barrier $\vartheta = 0.1$. This is clearly depicted in Figure 3(c) and (d) where WLCUSUM is antagonistic to GLR for both lengths $K = 5, 10$. Of course we must not forget that this performance of WLCUSUM is enjoyed at a significantly lower computational complexity level as compared to the window-limited GLR.

C. Different Parametric Families

We also consider the case with non-Gaussian pre-change distribution and Gaussian post-change distributions, i.e., the pre- and post-change does not belong to the same parametric family. In this case, the post-change parameter set Θ can be set as large as the entire space \mathbb{R}^K . We assume the pre-change distribution is a Laplace distribution with mean zero and variance 1, i.e., the density function $f_\infty = \frac{1}{\sqrt{2}} \exp\{-\sqrt{2}|x|\}$; and the post-change distribution is a normal distribution $\mathcal{N}(\mu, \sigma^2)$. Therefore, the pre-change distribution does not belong to the

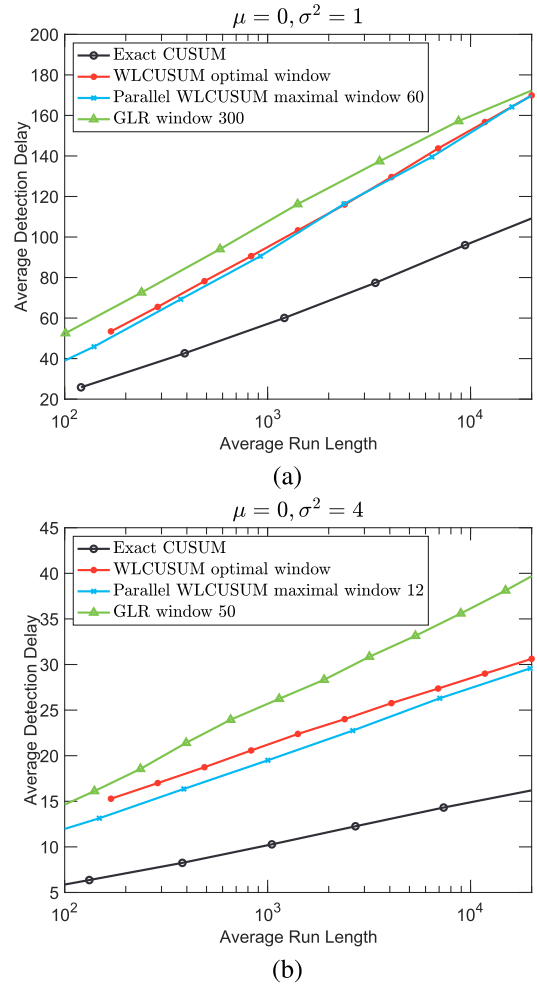


Fig. 4. Detection delays of different methods for univariate shift from Laplace to Normal distributions with (a) $\mu = 0, \sigma^2 = 1$ and assume known σ^2 ; (b) $\mu = 0, \sigma^2 = 4$ and assume unknown σ^2 .

post-change distribution family for any μ and σ , and we can safely choose the parameter set Θ as the entire parameter space.

In Figure 4(a) and (b) we performed experiments under two different detection scenarios. In particular in (a) we consider the post-change distribution being $\mathcal{N}(0, 1)$, i.e., although the distribution shifts from Laplace to Normal, their mean and variance remain the same. We assume the post-change variance σ^2 is known and only estimate the mean parameter during the detection procedure. Furthermore, in (b) we consider the post-change distribution being $\mathcal{N}(0, 4)$ and assume σ^2 is unknown, so we have to estimate the post-change mean and variance jointly through maximum likelihood estimate. In such case, we have the parameter space $\Theta = \{-\infty < \mu < \infty, \sigma^2 > 0\}$. Again both versions of WLCUSUM have comparable performance. The window-limited GLR, despite its high complexity as a result of the corresponding large windows (especially for (a) where the change is challenging to detect), does not enjoy better performance than WLCUSUM.

VIII. CONCLUSION

In this work we consider the sequential change detection problem with known pre-change distribution and unknown

post-change distribution but in certain parametric forms. We propose a window limited CUSUM procedure that uses a sliding window to perform online estimate for the unknown post-change parameter. A careful analysis on the average run length and detection delay shows the asymptotic optimality of the proposed method, with a window size much smaller than that required by window-limited GLR approach. The proposed framework also opens opportunities for several future research directions. For example, we may consider more efficient (such as “one-sample update” via stochastic gradient descent) estimate of the post-change parameter. Moreover, we may also try to relax the current constraint that the pre-change distribution must be a known distance away from the post-change distribution sets. The extension to joint detection and estimation is also worthwhile investigating in the future.

APPENDIX

Proof of Theorem 1: Recall the definition of the process $\{\mathcal{U}_t\}$ is $\mathcal{U}_t = \mathcal{U}_{t-1} + \log \frac{f_0(\xi_t, \hat{\theta}_{t-1})}{f_\infty(\xi_t)}$ with $\mathcal{U}_w = 0$, and the corresponding stopping time $\mathcal{T}' = \inf\{t > w : \mathcal{U}_t \geq \nu\}$. For the expectation $E_0^\theta[\mathcal{U}_{\mathcal{T}'}]$ we cannot apply the usual form of Wald’s identity because the terms under the sum are w -dependent. Indeed, while $\log \frac{f_0(\xi_t, \theta)}{f_\infty(\xi_t)}$ is independent from \mathcal{F}_{t-1} , when we replace the unknown θ with the estimate $\hat{\theta}_{t-1}$ then $\log \frac{f_0(\xi_t, \hat{\theta}_{t-1})}{f_\infty(\xi_t)}$ depends on \mathcal{F}_{t-1} and is independent only from \mathcal{F}_{t-w-1} due to the w -dependency of the estimates $\{\hat{\theta}_t\}$. There exists version of Wald’s identity for dependent samples [52] and for our analysis we are going to borrow ideas from this work but we intend to present all the details for our particular case.

For simplicity let us denote with $\ell_t = \log \frac{f_0(\xi_t, \hat{\theta}_{t-1})}{f_\infty(\xi_t)}$, then we observe that for $t > w$ we can write

$$\begin{aligned} E_0^\theta[\mathcal{U}_{\mathcal{T}'}] &= E_0^\theta \left[\sum_{t=w+1}^{\mathcal{T}'} \ell_t \right] \\ &= E_0^\theta \left[\sum_{t=w+1}^{\mathcal{T}'+w} \ell_t \right] - E_0^\theta \left[\sum_{t=\mathcal{T}'+1}^{\mathcal{T}'+w} \ell_t \right]. \end{aligned} \quad (20)$$

Consider the two terms in (20) separately. For the first we have

$$\begin{aligned} E_0^\theta \left[\sum_{t=w+1}^{\mathcal{T}'+w} \ell_t \right] &= E_0^\theta \left[\sum_{t=w+1}^{\infty} \ell_t \mathbb{1}_{\{\mathcal{T}' \geq t-w\}} \right] \\ &= E_0^\theta \left[\sum_{t=w+1}^{\infty} E_0^\theta[\ell_t | \mathcal{F}_{t-w-1}] \mathbb{1}_{\{\mathcal{T}' \geq t-w\}} \right] \\ &= E_0^\theta \left[\sum_{t=w+1}^{\infty} E_0^\theta[\ell_t] \mathbb{1}_{\{\mathcal{T}' \geq t-w\}} \right] \\ &= E_0^\theta[\ell_{w+1}] E_0^\theta[\mathcal{T}'] = \hat{l}_0 E_0^\theta[\mathcal{T}'], \end{aligned}$$

with the second equality being true since $\mathbb{1}_{\{\mathcal{T}' \geq t-w\}}$ is \mathcal{F}_{t-w-1} -measurable and the third equality being valid because $\ell_t = \log \frac{f_0(\xi_t, \hat{\theta}_{t-1})}{f_\infty(\xi_t)}$ is independent from \mathcal{F}_{t-w-1} .

Consider now the second term in (20), we observe that

$$\begin{aligned} E_0^\theta \left[\sum_{t=\mathcal{T}'+1}^{\mathcal{T}'+w} \ell_t \right] &= E_0^\theta \left[\sum_{t=w+1}^{\infty} \ell_t \mathbb{1}_{\{t > \mathcal{T}'\}} \mathbb{1}_{\{\mathcal{T}' \geq t-w\}} \right] \\ &= E_0^\theta \left[\sum_{t=w+1}^{\infty} E_0^\theta[\ell_t | \mathcal{F}_{t-1}] \mathbb{1}_{\{t > \mathcal{T}'\}} \mathbb{1}_{\{\mathcal{T}' \geq t-w\}} \right] \\ &\leq E_0^\theta \left[\sum_{t=w+1}^{\infty} l_0 \mathbb{1}_{\{t > \mathcal{T}'\}} \mathbb{1}_{\{\mathcal{T}' \geq t-w\}} \right] = w l_0, \end{aligned} \quad (21)$$

where for the inequality we used (12). Combining the two expressions we conclude that

$$\begin{aligned} \hat{l}_0 E_0^\theta[\mathcal{T}'] - w l_0 &\leq E_0^\theta[\mathcal{U}_{\mathcal{T}'}] \\ \Rightarrow E_0^\theta[\mathcal{T}'] &\leq \frac{E_0^\theta[\mathcal{U}_{\mathcal{T}'}] + w l_0}{\hat{l}_0} = \frac{E_0^\theta[\mathcal{U}_{\mathcal{T}'} - \nu] + \nu + w l_0}{\hat{l}_0}. \end{aligned} \quad (22)$$

The next step involves the control of the expectation of the overshoot $\mathcal{R}_\nu = \mathcal{U}_{\mathcal{T}'} - \nu$. In the case where \mathcal{U}_t is a sum of i.i.d. terms, we have from [51] an elegant result that bounds the average overshoot uniformly over all ν by a constant. Unfortunately, we were not able to produce a similar conclusion for the w -dependent case. Instead we developed an upper bound that increases as $\sqrt{\nu}$. The good news is that even with this cruder bound, the asymptotic characteristics of our scheme will turn out to be of the same order as the ones we would have enjoyed with a constant upper bound applied on the average overshoot.

We borrow ideas from [51] and modify them to accommodate the w -dependency. For any threshold $x > 0$ define \mathcal{T}'_x to be the corresponding stopping time

$$\mathcal{T}'_x = \inf\{t > w : \mathcal{U}_t \geq x\},$$

and denote the overshoot function as $\mathcal{R}_x = \mathcal{U}_{\mathcal{T}'_x} - x$. Define a sequence of stopping times $\{\tau_j\}$ with $\tau_0 = w$ and

$$\tau_j = \inf \left\{ t > \tau_{j-1} : \sum_{s=\tau_{j-1}+1}^t \ell_s > 0 \right\}, \quad j \geq 1,$$

and the corresponding ladder variables $z_j = \sum_{s=\tau_{j-1}+1}^{\tau_j} \ell_s > 0$. Due to the positivity of $\hat{l}_0 = E_0^\theta[\ell_t]$ we have that under the P_0^θ measure the stopping times are all a.s. finite. We can now see that $\mathcal{U}_{\tau_j} = \sum_{i=1}^j z_i$, in fact \mathcal{U}_t increases only at the stopping times $\{\tau_j\}$. For any given threshold ν , in order to stop at \mathcal{T}'_ν the statistic \mathcal{U}_t needs an increase at \mathcal{T}'_ν , which means that there exists a random index j_ν such that $\tau_{j_\nu} = \mathcal{T}'_\nu$. Due to this fact we can write $\mathcal{R}_\nu = \sum_{j=1}^{j_\nu} z_j - \nu$. Following the same steps as in [51], we observe that \mathcal{R}_x is a piecewise linear function and all pieces having slope -1 , thus we have

$$\begin{aligned} \int_0^\nu \mathcal{R}_x dx &= \int_0^{\mathcal{U}_{\mathcal{T}'_\nu}} \mathcal{R}_x dx - \int_\nu^{\mathcal{U}_{\mathcal{T}'_\nu}} \mathcal{R}_x dx \\ &= \frac{1}{2} \left\{ \sum_{j=1}^{j_\nu} z_j^2 - \mathcal{R}_\nu^2 \right\}. \end{aligned} \quad (23)$$

By definition $0 < z_j$ and $z_j = \sum_{t=\tau_{j-1}+1}^{\tau_j-1} \ell_t + \ell_{\tau_j}$ with $\sum_{t=\tau_{j-1}+1}^{\tau_j-1} \ell_t \leq 0$, then we have $0 < \ell_{\tau_j}$ and $0 < z_j \leq \ell_{\tau_j}$, which implies $z_j^2 \leq \ell_{\tau_j}^2 \leq \sum_{t=\tau_{j-1}+1}^{\tau_j} \ell_t^2$. Substituting in (23) yields

$$\begin{aligned} 2 \int_0^\nu \mathcal{R}_x dx &\leq \sum_{j=1}^{\hat{j}_\nu} \sum_{t=\tau_{j-1}+1}^{\tau_j} \ell_t^2 - \mathcal{R}_\nu^2 \\ &= \sum_{t=w+1}^{\mathcal{T}'} \ell_t^2 - \mathcal{R}_\nu^2 \leq \sum_{t=w+1}^{\mathcal{T}'+w} \ell_t^2 - \mathcal{R}_\nu^2. \end{aligned}$$

If we take the expectation of the previous expression and use Jensen's inequality on the last term we obtain

$$\begin{aligned} 0 &\leq 2 \int_0^\nu \mathbb{E}_0^\theta[\mathcal{R}_x] dx \leq \mathbb{E}_0^\theta \left[\sum_{t=w+1}^{\mathcal{T}'+w} \ell_t^2 \right] - (\mathbb{E}_0^\theta[\mathcal{R}_\nu])^2 \\ &= \frac{\hat{j}_0}{\hat{l}_0} \mathbb{E}_0^\theta \left[\sum_{t=w+1}^{\mathcal{T}'+w} \ell_t \right] - (\mathbb{E}_0^\theta[\mathcal{R}_\nu])^2 \\ &= \frac{\hat{j}_0}{\hat{l}_0} \mathbb{E}_0^\theta \left[\sum_{t=w+1}^{\mathcal{T}'} \ell_t + \sum_{t=\mathcal{T}'+1}^{\mathcal{T}'+w} \ell_t \right] - (\mathbb{E}_0^\theta[\mathcal{R}_\nu])^2 \\ &= \frac{\hat{j}_0}{\hat{l}_0} \mathbb{E}_0^\theta \left[\mathcal{U}_{\mathcal{T}'} + \sum_{t=\mathcal{T}'+1}^{\mathcal{T}'+w} \ell_t \right] - (\mathbb{E}_0^\theta[\mathcal{R}_\nu])^2 \\ &= \frac{\hat{j}_0}{\hat{l}_0} \mathbb{E}_0^\theta \left[\mathcal{R}_\nu + \nu + \sum_{t=\mathcal{T}'+1}^{\mathcal{T}'+w} \ell_t \right] - (\mathbb{E}_0^\theta[\mathcal{R}_\nu])^2 \\ &= \frac{\hat{j}_0}{\hat{l}_0} \left\{ \mathbb{E}_0^\theta[\mathcal{R}_\nu] + \nu + \mathbb{E}_0^\theta \left[\sum_{t=\mathcal{T}'+1}^{\mathcal{T}'+w} \ell_t \right] \right\} - (\mathbb{E}_0^\theta[\mathcal{R}_\nu])^2 \\ &\leq \frac{\hat{j}_0}{\hat{l}_0} \{ \mathbb{E}_0^\theta[\mathcal{R}_\nu] + \nu + w l_0 \} - (\mathbb{E}_0^\theta[\mathcal{R}_\nu])^2. \end{aligned}$$

The first equality is true because $\mathbb{E}_0^\theta[\sum_{t=w+1}^{\mathcal{T}'+w} \ell_t^2] = \hat{j}_0 \mathbb{E}_0^\theta[\mathcal{T}']$ and $\mathbb{E}_0^\theta[\sum_{t=w+1}^{\mathcal{T}'+w} \ell_t] = \hat{l}_0 \mathbb{E}_0^\theta[\mathcal{T}']$. Also for the last inequality we used (21). From the nonnegativity of the integral we have $(\mathbb{E}_0^\theta[\mathcal{R}_\nu])^2 \leq \frac{\hat{j}_0}{\hat{l}_0} \{ \mathbb{E}_0^\theta[\mathcal{R}_\nu] + \nu + w l_0 \}$, from which we conclude that $\mathbb{E}_0^\theta[\mathcal{R}_\nu] \leq \frac{\hat{j}_0}{\hat{l}_0} + \left(\frac{\hat{j}_0}{\hat{l}_0} \nu \right)^{\frac{1}{2}} + \left(\frac{\hat{j}_0}{\hat{l}_0} w l_0 \right)^{\frac{1}{2}}$. Given also from (14) that $\nu = \log \gamma$, substituting in (22) produces the desired upper bound. \square

REFERENCES

- [1] H. V. Poor and O. Hadjiladis, *Quickest Detection*. Cambridge, U.K.: Cambridge Univ. Press, 2008.
- [2] D. Siegmund, *Sequential Analysis: Tests Confidence Intervals* (Springer Series in Statistics). New York, NY, USA: Springer-Verlag, 1985.
- [3] A. Tartakovsky, I. Nikiforov, and M. Basseville, *Sequential Analysis: Hypothesis Testing and Changepoint Detection* (Monographs on Statistics and Applied Probability). London, U.K.: CRC Press, 2015.
- [4] L. Xie, S. Zou, Y. Xie, and V. V. Veeravalli, "Sequential (quickest) change detection: Classical results and new directions," *IEEE J. Sel. Areas Inf. Theory*, vol. 2, no. 2, pp. 494–514, Jun. 2021.
- [5] L. Xie, Y. Xie, and G. V. Moustakides, "Asynchronous multi-sensor change-point detection for seismic tremors," in *Proc. IEEE Int. Symp. Inf. Theory (ISIT)*, Jul. 2019, pp. 787–791.
- [6] J. Shi and S. Zhou, "Quality control and improvement for multistage systems: A survey," *IIE Trans.*, vol. 41, no. 9, pp. 744–753, Jul. 2009.
- [7] T. L. Lai, "Sequential changepoint detection in quality control and dynamical systems," *J. Statist. Soc. Ser. B*, vol. 57, pp. 613–658, Sep. 1995.
- [8] D. Balageas, C.-P. Fritzen, and A. Güemes, *Structural Health Monitoring*, vol. 90. Hoboken, NJ, USA: Wiley, 2010.
- [9] S. Li, Y. Xie, M. Farajtabar, A. Verma, and L. Song, "Detecting changes in dynamic events over networks," *IEEE Trans. Signal Inf. Process. Over Netw.*, vol. 3, no. 2, pp. 346–359, Jun. 2017.
- [10] V. Chandola, A. Banerjee, and V. Kumar, "Anomaly detection: A survey," *ACM Comput. Surv.*, vol. 41, no. 3, pp. 1–58, Jul. 2009.
- [11] A. G. Tartakovsky, "Rapid detection of attacks in computer networks by quickest changepoint detection methods," in *Data Analysis for Network Cyber-Security*. London, U.K.: Imperial College Press, 2014, pp. 33–70.
- [12] E. S. Page, "Continuous inspection schemes," *Biometrika*, vol. 41, nos. 1–2, pp. 100–115, 1954.
- [13] G. V. Moustakides, "Optimal stopping times for detecting changes in distributions," *Ann. Statist.*, vol. 14, no. 4, pp. 1379–1387, Dec. 1986.
- [14] T. Leung Lai, "Information bounds and quick detection of parameter changes in stochastic systems," *IEEE Trans. Inf. Theory*, vol. 44, no. 7, pp. 2917–2929, May 1998.
- [15] J. Chen, Y. Zhao, A. Goldsmith, and H. V. Poor, "Optimal joint detection and estimation in linear models," in *Proc. 52nd IEEE Conf. Decis. Control*, Dec. 2013, pp. 4416–4421.
- [16] G. V. Moustakides, G. H. Jajamovich, A. Tajer, and X. Wang, "Joint detection and estimation: Optimum tests and applications," *IEEE Trans. Inf. Theory*, vol. 58, no. 7, pp. 4215–4229, Jul. 2012.
- [17] G. Lorden, "Procedures for reacting to a change in distribution," *Ann. Math. Statist.*, vol. 42, no. 6, pp. 1897–1908, Dec. 1971.
- [18] M. Basseville and I. V. Nikiforov, *Detection Abrupt Changes: Theory Application* (Prentice-Hall Information and System Sciences Series). Englewood Cliffs, NJ, USA: Prentice-Hall, 1993.
- [19] A. G. Tartakovsky, *Sequential Change Detection Hypothesis Testing: Gen. Non-IID Stochastic Models Asymptotically Optimal Rules* (Monographs on Statistics and Applied Probability). New York, NY, USA: CRC Press, 2020.
- [20] V. V. Veeravalli and T. Banerjee, "Quickest change detection," in *Academic Press Library in Signal Processing: Array and Statistical Signal Processing*, vol. 3. Cham, Switzerland: Springer, 2014, pp. 209–255.
- [21] G. V. Moustakides, "Sequential change detection revisited," *Ann. Statist.*, vol. 36, no. 2, pp. 787–807, Apr. 2008.
- [22] A. G. Tartakovsky and G. V. Moustakides, "State-of-the-art in Bayesian changepoint detection," *Sequential Anal.*, vol. 29, no. 2, pp. 125–145, Apr. 2010.
- [23] A. N. Shiryaev, "On optimum methods in quickest detection problems," *Theory Probab. Appl.*, vol. 8, no. 1, pp. 22–46, Jan. 1963.
- [24] M. Pollak, "Optimal detection of a change in distribution," *Ann. Statist.*, vol. 13, no. 1, pp. 206–227, Mar. 1985.
- [25] A. S. Polunchenko and A. G. Tartakovsky, "On optimality of the Shiryaev–Roberts procedure for detecting a change in distribution," *Ann. Statist.*, vol. 38, no. 6, pp. 3445–3457, Dec. 2010.
- [26] C. Kirch and J. Tadjuidje Kamgaing, "On the use of estimating functions in monitoring time series for change points," *J. Stat. Planning Inference*, vol. 161, pp. 25–49, Jun. 2015.
- [27] J. M. Lucas, "Combined shewhart-CUSUM quality control schemes," *J. Quality Technol.*, vol. 14, no. 2, pp. 51–59, Apr. 1982.
- [28] J. O. Westgard, T. Groth, T. Aronsson, and C. H. de Verdier, "Combined Shewhart-CUSUM control chart for improved quality control in clinical chemistry," *Clin. Chem.*, vol. 23, no. 10, pp. 1881–1887, Oct. 1977.
- [29] R. S. Sparks, "CUSUM charts for signalling varying location shifts," *J. Quality Technol.*, vol. 32, no. 2, pp. 157–171, Apr. 2000.
- [30] Y. Zhao, F. Tsung, and Z. Wang, "Dual CUSUM control schemes for detecting a range of mean shifts," *IIE Trans.*, vol. 37, no. 11, pp. 1047–1057, Nov. 2005.
- [31] Y. Yu, O. H. M. Padilla, D. Wang, and A. Rinaldo, "A note on online change point detection," 2020, [arXiv:2006.03283](https://arxiv.org/abs/2006.03283).
- [32] G. Romano, I. Eckley, P. Fearnhead, and G. Rigall, "Fast online changepoint detection via functional pruning CUSUM statistics," 2021, [arXiv:2110.08205](https://arxiv.org/abs/2110.08205).
- [33] I. V. Nikiforov, "A generalized change detection problem," *IEEE Trans. Inf. Theory*, vol. 41, no. 1, pp. 171–187, 1995.
- [34] T. Leung Lai, "Sequential multiple hypothesis testing and efficient fault detection-isolation in stochastic systems," *IEEE Trans. Inf. Theory*, vol. 46, no. 2, pp. 595–608, Mar. 2000.

- [35] S. Abbasi and A. Haq, "Optimal CUSUM and adaptive CUSUM charts with auxiliary information for process mean," *J. Stat. Comput. Simul.*, vol. 89, no. 2, pp. 337–361, Jan. 2019.
- [36] S. Abbasi and A. Haq, "New adaptive CUSUM charts for process mean," *Commun. Statist.-Simul. Comput.*, vol. 49, no. 11, pp. 2944–2962, Nov. 2020.
- [37] W. Jiang, L. Shu, and D. W. Apley, "Adaptive CUSUM procedures with EWMA-based shift estimators," *IIE Trans.*, vol. 40, no. 10, pp. 992–1003, Aug. 2008.
- [38] Y. Luo, Z. Li, and Z. Wang, "Adaptive CUSUM control chart with variable sampling intervals," *Comput. Statist. Data Anal.*, vol. 53, no. 7, pp. 2693–2701, May 2009.
- [39] L. Shu and W. Jiang, "A Markov chain model for the adaptive CUSUM control chart," *J. Quality Technol.*, vol. 38, no. 2, pp. 135–147, Apr. 2006.
- [40] Z. Wu, J. Jiao, M. Yang, Y. Liu, and Z. Wang, "An enhanced adaptive CUSUM control chart," *IIE Trans.*, vol. 41, no. 7, pp. 642–653, May 2009.
- [41] Y. Cao, L. Xie, Y. Xie, and H. Xu, "Sequential change-point detection via online convex optimization," *Entropy*, vol. 20, no. 2, p. 108, Feb. 2018.
- [42] G. Lorden and M. Pollak, "Sequential change-point detection procedures that are nearly optimal and computationally simple," *Sequential Anal.*, vol. 27, no. 4, pp. 476–512, Nov. 2008.
- [43] Y. Wu, "Detecting changes in a multiparameter exponential family by using adaptive CUSUM procedure," *Sequential Anal.*, vol. 36, no. 4, pp. 467–480, Oct. 2017.
- [44] Q. Xu and Y. Mei, "Multi-stream quickest detection with unknown post-change parameters under sampling control," in *Proc. IEEE Int. Symp. Inf. Theory (ISIT)*, Jul. 2021, pp. 112–117.
- [45] B. W. Silverman, *Density Estimation for Statistics and Data Analysis* (Monographs on Statistics and Applied Probability). London, U.K.: Chapman & Hall, 1986.
- [46] Q. Xu, Y. Mei, and G. V. Moustakides, "Optimum multi-stream sequential change-point detection with sampling control," *IEEE Trans. Inf. Theory*, vol. 67, no. 11, pp. 7627–7636, Nov. 2021.
- [47] S. Janson, "Runs in m -dependent sequences," *Ann. Probab.*, vol. 12, no. 3, pp. 805–818, 1984.
- [48] L. Xie, Y. Xie, and G. V. Moustakides, "Sequential subspace change point detection," *Sequential Anal.*, vol. 39, no. 3, pp. 307–335, Jul. 2020.
- [49] T. Jacob and R. K. Bansal, "Sequential change detection based on universal compression algorithms," in *Proc. IEEE Int. Symp. Inf. Theory*, Jul. 2008, pp. 1108–1112.
- [50] E. L. Lehmann and G. Casella, *Theory Point Estimation*. Cham, Switzerland: Springer, 2006.
- [51] G. Lorden, "On excess over the boundary," *Ann. Math. Statist.*, vol. 41, no. 2, pp. 520–527, Apr. 1970.
- [52] G. V. Moustakides, "Extension of Wald's first lemma to Markov processes," *J. Appl. Probab.*, vol. 36, no. 1, pp. 48–59, Mar. 1999.

Liyan Xie received the B.Sc. degree in statistics from the University of Science and Technology of China in 2016 and the Ph.D. degree in industrial engineering from the Georgia Institute of Technology in 2021. She is currently an Assistant Professor with the School of Data Science, The Chinese University of Hong Kong, Shenzhen. Her research interests include the intersection between statistics and optimization, with a primary focus on sequential change detection and their applications in sensor networks and healthcare. She was a recipient of the 2020 ARC Student Fellowship and the IDEaS-TRIAD Research Scholarships from the Transdisciplinary Research Institute for Advancing Data Science (TRIAD) at Georgia Tech. She was awarded the Finalist of the INFORMS QSR Student Paper Competition and the Runner-Up for the INFORMS Computing Society Student Paper Prize in 2019.

George V. Moustakides (Life Senior Member, IEEE) received the Diploma degree in electrical and mechanical engineering from the National Technical University of Athens, Greece, in 1979, the M.Sc. degree in systems engineering from the Moore School of Electrical Engineering, University of Pennsylvania, Philadelphia, in 1980, and the Ph.D. degree in electrical engineering and computer science from Princeton University, Princeton, NJ, USA, in 1983. In 2007, he joined the Department of Electrical and Computer Engineering, University of Patras, Patras, Greece, where he is currently an Emeritus Professor. He held various long-term positions in the past: a Junior and Senior Researcher with INRIA, France; a Professor with the University of Thessaly, Greece; and a long-term Visiting Professor with Rutgers University. He has also been a Visiting Scholar and an Adjunct Professor with Princeton University, the University of Pennsylvania, Columbia University, the University of Maryland, the Georgia Institute of Technology, the University of Southern California, the University of Illinois at Urbana-Champaign, Rutgers University, and Aalto University, Finland. His research interests include sequential detection, statistical signal processing, and statistical machine learning. From 2011 to 2014, he served as an Associate Editor for *Detection and Estimation Theory* and an (inaugural) Associate Editor for Sequential Methods for the IEEE TRANSACTIONS ON INFORMATION THEORY from 2016 to 2018.

Yao Xie (Member, IEEE) received the Ph.D. degree in electrical engineering (mathematics) from Stanford University. She was a Research Scientist with Duke University. She is the Coca-Cola Foundation Chair and Professor at the Georgia Institute of Technology in the H. Milton Stewart School of Industrial and Systems Engineering and an Associate Director of the Machine Learning Center. Her research interests include the intersection of statistics, machine learning, and optimization in providing theoretical guarantees and developing computationally efficient and statistically powerful methods for problems motivated by real-world applications. She received the National Science Foundation (NSF) CAREER Award in 2017, the INFORMS Wagner Prize Finalist in 2021, and the INFORMS Gaver Early Career Award for Excellence in Operations Research in 2022. She is the Area Chair of NeurIPS and ICML. She is currently an Associate Editor of IEEE TRANSACTIONS ON INFORMATION THEORY, IEEE TRANSACTIONS ON SIGNAL PROCESSING, *Journal of the American Statistical Association*, *Theory and Methods, Sequential Analysis: Design Methods and Applications*, and *INFORMS Journal on Data Science*.