A binning approach to quickest change detection with unknown post-change distribution

Lau, Tze Siong; Tay, Wee Peng; Veeravalli, Venugopal V.

2018

Lau, T. S., Tay, W. P., & Veeravalli, V. V. (2019). A binning approach to quickest change detection with unknown post-change distribution. *IEEE Transactions on Signal Processing*, 67(3), 609-621. doi:10.1109/TSP.2018.2881666

<https://hdl.handle.net/10356/137146>

<https://doi.org/10.1109/TSP.2018.2881666>

© 2018 IEEE. Personal use of this material is permitted. Permission from IEEE must be obtained for all other uses, in any current or future media, including reprinting/republishing this material for advertising or promotional purposes, creating new collective works, for resale or redistribution to servers or lists, or reuse of any copyrighted component of this work in other works. The published version is available at: <https://doi.org/10.1109/TSP.2018.2881666>.

Downloaded on 26 Apr 2023 19:20:32 SGT

A Binning Approach to Quickest Change Detection with Unknown Post-Change Distribution

Tze Siong Lau, Wee Peng Tay, *Senior Member, IEEE*, and Venugopal V. Veeravalli, *Fellow, IEEE*

Abstract—The problem of quickest detection of a change in distribution is considered under the assumption that the pre-change distribution is known, and the post-change distribution is only known to belong to a family of distributions distinguishable from a discretized version of the pre-change distribution. A sequential change detection procedure is proposed that partitions the sample space into a finite number of bins, and monitors the number of samples falling into each of these bins to detect the change. A test statistic that approximates the generalized likelihood ratio test is developed. It is shown that the proposed test statistic can be efficiently computed using a recursive update scheme, and a procedure for choosing the number of bins in the scheme is provided. Various asymptotic properties of the test statistic are derived to offer insights into its performance trade-off between average detection delay and average run length to false alarm. Testing on synthetic and real data demonstrates that our approach is comparable or better in performance to existing non-parametric change detection methods.

Index Terms—Quickest change detection, non-parametric test, Generalized Likelihood Ratio Test (GLRT), average run length, average detection delay

I. INTRODUCTION

Quickest change detection (QCD) is a fundamental problem in statistics. Given a sequence of observations that have a certain distribution up to an unknown change point ν , and have a different distribution after that, the goal is to detect this change in distribution as quickly as possible subject to false alarm constraints. The QCD problem arises in many practical situations, including applications in manufacturing such as quality control, where any deviation in the quality of products must be quickly detected. With the increase in the amount and types of data modern-day sensors are able to observe, sequential change detection methods have also found applications in the areas of power system line outage detection [1], [2], bioinformatics [3], network surveillance [4]–[9], fraud detection [10], structural health monitoring [11], spam detection [12], spectrum reuse [13]–[19], video segmentation [20], and resource allocation and scheduling [21], [22]. In many of these applications, the detection algorithm has to operate in real time with reasonable computation complexity.

This work was supported in part by the Singapore Ministry of Education Academic Research Fund Tier 2 grant MOE2014-T2-1-028, by the US National Science Foundation under grant CCF 15-14245, and by the Air Force Office of Scientific Research (AFOSR) under grant FA9550-16-1-0077.

T. S. Lau and W. P. Tay are with the School of Electrical and Electronic Engineering, Nanyang Technological University, Singapore (e-mail: TLAU001@e.ntu.edu.sg, wptay@ntu.edu.sg). V. V. Veeravalli is with the Coordinated Science Laboratory and the Department of Electrical and Computer Engineering, University of Illinois at Urbana-Champaign, Urbana, IL, USA (e-mail: vvv@illinois.edu).

In this paper, we consider the QCD problem with a known pre-change distribution and unknown post-change distribution. The only information we have about the post-change distribution is that it belongs to a set of distributions that is distinguishable from the pre-change distribution when the sample space is discretized into N bins. Assuming that the pre-change distribution is known is reasonable, because in most practical applications, a large amount of data generated by the pre-change distribution is available to the observer who may use this data to obtain an accurate approximation of the pre-change distribution [23], [24]. However, estimating or even modelling the post-change distribution is often impractical as we may not know *a priori* what kind of change will happen. We seek to design a low-complexity detection algorithm that allows us to quickly detect the change, under false alarm constraints, and with minimal knowledge of the post-change distribution. To solve this problem, we propose a new test statistic based on binning and the generalized likelihood ratio test (GLRT), which approximates the Cumulative Sum (CuSum) statistic of Page [25]. We propose a method to choose the appropriate number of bins, and show that our proposed test statistic can be updated sequentially in a manner similar to the CuSum statistic. We also provided an analysis of the performance of the proposed test statistic.

A. Related Works

For the case where the pre- and post-change distributions are known, Page [25] developed the CuSum Control Chart for QCD, which was later shown to be asymptotically optimal under a certain criterion by Lorden [26], and exactly optimal under Lorden's criterion by Moutakides [27]. We refer the reader to [28]–[30] and the references therein for an overview of the QCD problem. There are also many existing works that consider the QCD problem where the post-change distribution is unknown to a certain degree. In [31], the authors considered the case where the post-change distribution belongs to a one-parameter exponential family with the pre-change distribution being known. The case where both the pre- and post-change distributions belong to a one-parameter exponential family was considered by [32]. In [33], the authors developed a data-efficient scheme that allows for optional sampling of the observations in the case when either the post-change family of distributions is finite, or both the pre- and post-change distributions belong to a one parameter exponential family. In [34], the authors assumed that the observations are time-dependent as in ARMA, general linear processes and α -mixing processes. In [35], the authors proposed using a

infinite hidden Markov model for tracking parametric changes in the signal model. Classical approaches to the QCD problem without strong distributional assumptions can be found in [36]–[39]. Although there are no distributional assumptions, the type of change assumed in [36]–[40] is a shift in the mean and in [41], a shift in the scale of the observations. In [42], the authors provided a kernel-based detection scheme for a change-point detection problem where the post-change distribution is completely unknown. However, the kernel-based detection scheme requires choosing a proper kernel bandwidth. This is usually done by tuning the bandwidth parameter on a reference dataset representative of the pre-change and possible post-change conditions. In [43], the authors proposed a non-parametric algorithm for detecting a change in mean. They assume that the mean of the pre- and post-change distributions are not equal but do not assume any knowledge of the distribution function.

In this paper, we consider the case where less information is available about the post-change distribution, compared to [31]–[33]. The only assumption we make is that the post-change distribution belongs to a set of distributions that are distinguishable from the pre-change distribution when the sample space is discretized into a known number of bins. Our method is also more general than [36]–[41], [43] as it is not restricted to a shift in mean, location or scale. Related work is discussed in [44], in which an asymptotically optimal universal scheme is developed to isolate an outlier data stream that experiences a change, from a pool of at least two typical streams. In our work, we only have a single stream of observations and therefore our scheme is unable to make use of comparisons across streams to detect for changes. Other related works include [37], [42], [45], which do not require prior information about the pre-change distribution. These works first partition the observations using a hypothesized change point into two sample sets and then test if they are generated by the same distribution using a two-sample test. This approach results in weak detection performances when the change occurs early in the observation sequence. In our work, we assume that prior information about the pre-change distribution is available, which allows us to detect the change even when the change occurs early in the observations.

B. Our Contributions

In this paper, we consider the QCD problem where the pre-change distribution is known, and the post-change distribution belongs to a family of distributions distinguishable from the pre-change distribution when the sample space is discretized into bins. Our goal is to develop an algorithm that is effective and has low computational complexity. Our main contributions are as follows:

- 1) We propose a Binned Generalized CuSum (BG-CuSum) test and derive a recursive update formula for our test statistic, which allows for an efficient implementation of the test.
- 2) We derive a lower bound for the average run length (ARL) to false alarm of our BG-CuSum test, and show asymptotic properties of the BG-CuSum statistic in order

to provide insights into the trade-off between the average detection delay (ADD) and ARL.

- 3) We provide simulations and experimental results, which indicate that our proposed BG-CuSum test outperforms various other non-parametric change detection methods and its performance approaches a test with known asymptotic properties as the ARL becomes large.

A preliminary version of this work was presented in [46]. To the best of our knowledge, the only other recursive test known in the literature for the case where the post-change distribution is not completely specified and does not belong to a finite set of distributions is discussed in [47]. The paper [47] derived a scheme that uses 3 registers for storing past information to detect a change in the parameter of an exponential family. Our work is applicable to more general changes in distribution as we do not assume that the post-change distribution belong to an exponential family. Due to technical difficulties, we are not able to obtain the asymptotic ADD for our BG-CuSum test. Instead, we derive the asymptotic ADD for a related non-recursive test, which is asymptotically optimal when the pre- and post-change distributions are discrete. Simulations indicate that the BG-CuSum test has similar asymptotic behavior as this latter test, which is however computationally expensive. The asymptotic performance analysis of the BG-CuSum test remains an open problem.

The rest of this paper is organized as follows. In Section II, we present the QCD signal model and problem formulation. In Section III, we present our GLRT based QCD procedure and propose an approximation that can be computed efficiently. In Section IV, we analyse the asymptotic behavior of our test statistic, and provide a heuristic analysis of the trade-off between the ADD and the ARL to false alarm. In Section V, we present simulation and experimental results to illustrate the performance of our algorithm. Finally, in Section VI, we present some concluding remarks.

Notations: We use $\mathbb{R}$ and $\mathbb{N}$ to denote the set of real numbers and positive integers, respectively. The operator $\mathbb{E}^f$ denotes mathematical expectation with respect to (w.r.t.) the probability distribution with generalized probability density (pdf) f , and $X \sim f$ means that the random variable X has distribution with pdf f . We use upper-case (e.g. X) to refer to a random variable, and lower case (e.g. x) to refer to a realization of the random variable X . We let $\mathbb{P}_\nu$ and $\mathbb{E}_\nu$ denote the measure and mathematical expectation respectively where the change point is at ν for $\nu \in \mathbb{N}$. In particular, $\mathbb{P}_1$ and $\mathbb{E}_1$ denote the measure and mathematical expectation, respectively, when the change has occurred at time $t=1$, i.e., all the observations are distributed according to the post-change distribution. Similarly, we let $\mathbb{P}_\infty$ and $\mathbb{E}_\infty$ denote the measure and mathematical expectation, respectively, when there is no change, i.e., all the observations are distributed according to the pre-change distribution. The Gaussian distribution with mean μ and variance σ^2 is denoted as $\mathcal{N}(\mu, \sigma^2)$. The Dirac delta function at θ is denoted as δ_θ . $|A|$ denotes the cardinality of the set A .

II. PROBLEM FORMULATION

We consider distributions defined on the sample space $\mathbb{R}$. Let $f = p_0 f_c + \sum_{h=1}^H p_h \delta_{\theta_h}$ be the generalized pdf of the pre-change distribution, where f_c is the pdf of the continuous part, $\{\theta_h \in \mathbb{R}: h=1, \dots, H\}$ are the locations of the point masses, and $0 \leq p_h \leq 1$ for all $h \geq 0$ with $\sum_{h=0}^H p_h = 1$. Similarly, let $g = q_0 g_c + \sum_{h=1}^H q_h \delta_{\theta_h}$ be the generalized pdf of the post-change distribution where $g \neq f$, and $0 \leq q_h \leq 1$ for all $h \geq 0$ with $\sum_{h=0}^H q_h = 1$. Let $X_1, X_2, \dots$ be a sequence of real valued random variables satisfying the following:

$$\begin{cases} X_t \sim f & \text{i.i.d. for all } t < \nu, \\ X_t \sim g & \text{i.i.d. for all } t \geq \nu, \end{cases} \quad (1)$$

where $\nu \geq 0$ is an unknown but deterministic change point. The quickest change detection problem is to detect the change in distribution, through observing $X_1 = x_1, X_2 = x_2, \dots$, as quickly as possible while keeping the false alarm rate low. In this paper, we assume that the observer only knows the pre-change distribution f but does not have full knowledge of the post-change distribution g . We assume that the post-change distribution is distinguishable from the pre-change distribution when sample space is discretized into a known number of bins, defined as follows.

Definition 1. Let $\Theta = \{\theta_1, \dots, \theta_H\}$. For any $N \geq 1$, let $I_1^N = (-\infty, z_1] \setminus \Theta$, $I_2^N = (z_1, z_2] \setminus \Theta$, ..., and $I_N^N = (z_{N-1}, \infty) \setminus \Theta$ be N sets such that for each $j \in \{1, \dots, N\}$, we have $\int_{I_j^N} f_c(x) dx = \frac{1}{N}$. Also define the sets $I_{N+1}^N = \{\theta_1\}, \dots, I_{N+H}^N = \{\theta_H\}$. We call each of the sets $I_j^N, j=1, \dots, N+H$, a bin.

Definition 2. A distribution g , absolutely continuous w.r.t. f , is distinguishable from f w.r.t. N if there exists $j \in \{1, \dots, N+H\}$ such that $\int_{I_j^N} f(x) dx \neq \int_{I_j^N} g(x) dx$. The set $D(f, N)$ is the family of all generalized pdfs g distinguishable from f w.r.t. N .

Any distribution $g \neq f$ is distinguishable from f for N sufficiently large. To see why this is true, we define the functions g_N and f_N as

$$g_N(x) \doteq \int_{I_j^N} g(y) dy, \quad f_N(x) \doteq \int_{I_j^N} f(y) dy, \quad (2)$$

where j is the unique integer such that $x \in I_j^N$. If there exists $h \in \{0, \dots, H\}$ such that $p_h \neq q_h$ then $g \in D(f, 1)$. Now suppose that $p_h = q_h$ for all $h \in \{0, \dots, H\}$. Let F_c and G_c be the cumulative distribution function of f_c and g_c , respectively. Since both F_c and G_c are continuous and $F_c \neq G_c$, there exists an interval J such that for any $x \in J$, $F_c(x) \neq G_c(x)$. Then for any $N > 1/\int_J f_c(x) dx$, there exists some x such that $g_N(x) \neq f_N(x)$ because otherwise $G_c(x) = F_c(x)$ for any x that is a boundary point of a set in $\{I_j^N: j=1, \dots, N\}$.

In the supplementary material [48], we give an exact characterization of N such that a distribution g is distinguishable from f w.r.t. N . We also provide an example to illustrate how N can be determined if additional moment information is available. For the rest of this paper, we make the following assumption.

Assumption 1. For a known positive integer N and pre-change distribution f , the post-change distribution g belongs to the set $D(f, N)$.

In a typical sequential change detection procedure, at each time t , a test statistic $S(t)$ is computed based on the currently available observations $X_1 = x_1, \dots, X_t = x_t$, and the observer makes the decision that a change has occurred at a stopping time τ , where τ is the first t such that $S(t)$ exceeds a pre-determined threshold b :

$$\tau(b) = \inf\{t: S(t) \geq b\}. \quad (3)$$

We evaluate the performance of a change detection scheme using the ARL to false alarm and the worst-case ADD (WADD) using Lorden's definitions [26] as follows:

$$\begin{aligned} \text{ARL}(\tau) &= \mathbb{E}_\infty[\tau], \\ \text{WADD}(\tau) &= \sup_{\nu \geq 1} \text{esssup} \mathbb{E}_\nu[(\tau - \nu + 1)^+ | X_1, \dots, X_{\nu-1}]. \end{aligned}$$

where $\sup$ is the supremum operator and esssup is the essential supremum operator. The quickest change detection problem can be formulated as the following minimax problem [26]:

$$\min_{\tau} \text{WADD}(\tau), \quad (4)$$

subject to $\text{ARL}(\tau) \geq \gamma$,

for some given $\gamma > 0$. We refer the interested reader to Chapter 6 of [28] for a comprehensive treatment and overview of Lorden's formulation of the quickest change detection problem.

III. TEST STATISTIC BASED ON BINNING

In this section, we derive a test statistic that can be recursively updated. If the post-change distribution belongs to a finite set of distributions, the GLRT statistic is the maximum of a finite set of CuSum statistics, one corresponding to each possible post-change distribution, and therefore has a recursion. To the best of our knowledge, the only other recursive test known in the literature for the case where the post-change distribution is not completely specified and does not belong to a finite set of distributions is proposed in [47]. The method in [47] requires that both the pre- and post-change distributions belong to a single parameter exponential family.

Suppose that the post-change distribution $g \in D(f, N)$ and we observe the sequence $X_1 = x_1, X_2 = x_2, \dots$. If g_N and f_N are both known, Page's CuSum test statistic [25] for the binned observations is

$$S(t) = \max_{1 \leq k \leq t+1} \sum_{j=k}^t \log \frac{g_N(x_j)}{f_N(x_j)}. \quad (5)$$

Note that $S(t)$ in (5) takes the value 0 if $k=t+1$ is the maximizer. The test statistic $S(t)$ has a convenient recursion $S(t+1) = \max\{S(t) + \log(g_N(x_{t+1})/f_N(x_{t+1})), 0\}$.

In our problem formulation, g_N is unknown. We thus replace $g_N(x_i)$ in (5) with its maximum likelihood estimator

$$g_N^{k:t}(x_i) = \frac{|\{x_r: k \leq r \leq t \text{ and } x_r \in I_j^N\}|}{t-k+1},$$

where j is the unique integer such that $x_i \in I_j^N$. Note that in computing $g_N^{k:t}(x_i)$, we use only the samples $x_k, \dots, x_t$. We then have the test statistic

$$\max_{1 \leq k \leq t+1} \sum_{i=k}^t \log \frac{g_N^{k:t}(x_i)}{f_N(x_i)}. \quad (6)$$

In the case where $t-k+1$ is small, the maximum likelihood estimator $g_N^{k:t}$ tends to over-fit the observed data and evaluating the likelihood of x_i using $g_N^{k:t}(x_i)$ biases the test statistic. Furthermore, this couples the estimation of g_N and the instantaneous likelihood ratio $\frac{g_N(x_i)}{f_N(x_i)}$. In order to compensate for this over-fitting, we choose not to include observations $x_i, \dots, x_t$ in the estimation of g_N . This also decouples the estimation of g_N and the likelihood ratio $\frac{g_N(x_i)}{f_N(x_i)}$. However, if x_i is the first observation occurring in the set I_j^N , we have $g_N^{k:i-1}(x_i) = 0$. To avoid this, we define the regularized version of $g_N^{k:i-1}$ as

$$\hat{g}_N^{k:i-1}(x_i) = \begin{cases} \frac{|\{x_r: k \leq r \leq i-1 \text{ and } x_r \in I_j^N\}| + R}{(N+H)R + i - k} & \text{if } k \leq i-1, \\ f_N(x_i) & \text{otherwise,} \end{cases} \quad (7)$$

where R is a fixed positive constant, and j is the unique integer such that $x_i \in I_j^N$. Our test statistic then becomes

$$\hat{S}_N(t) = \max_{1 \leq k \leq t+1} \sum_{i=k}^t \log \frac{\hat{g}_N^{k:i-1}(x_i)}{f_N(x_i)}, \quad (8)$$

with the stopping time

$$\hat{\tau}(b) = \inf\{t: \hat{S}_N(t) > b\}. \quad (9)$$

In practice, R is chosen to be of the order of N so that $\hat{g}_N^{k:i-1}(x_i)$ approaches $1/N$ as $N \rightarrow \infty$. This controls the variability of (8) by controlling the range of values that $\log \frac{\hat{g}_N^{k:i-1}(x_i)}{f_N(x_i)}$ can take. Computation of the test statistic (8) is inefficient as the estimator $\hat{g}_N^{k:i-1}$ needs to be recomputed each time a new observation $X_t = x_t$ is made, leading to computational complexity increasing linearly w.r.t. t . One way to prevent this increase in computational complexity is by searching for a change point from the previous most likely change point rather than from $t=1$, and also using observations from the previous most likely change point to the current observation to update the estimator for g . Our proposed BG-CuSum test statistic $\tilde{S}_N$ and test $\tilde{\tau}$ are defined as follows: For each $t \geq 1$,

$$\tilde{S}_N(t) = \max_{\lambda_{t-1} \leq k \leq t+1} \sum_{i=k}^t \log \frac{\hat{g}_N^{\lambda_{t-1}:i-1}(x_i)}{f_N(x_i)}, \quad (10)$$

$$\lambda_t = \max \left\{ \underset{\substack{\lambda_{t-1} \leq k \leq t+1 \\ k \neq \lambda_{t-1}+1}}{\operatorname{argmax}} \sum_{i=k}^t \log \frac{\hat{g}_N^{\lambda_{t-1}:i-1}(x_i)}{f_N(x_i)} \right\}, \quad (11)$$

$$\tilde{\tau}(b) = \inf\{t: \tilde{S}_N(t) \geq b\}, \quad (12)$$

where $\lambda_0 = 1$, and b is a fixed threshold. The outer maximum in (11) is to ensure that λ_t is uniquely defined when there is more than one maximizer. Due to the design of the estimator (7), we have $\hat{g}_N^{\lambda_{t-1}: \lambda_{t-1}-1}(x_{\lambda_{t-1}}) = f_N(x_{\lambda_{t-1}})$ and

$$\sum_{i=\lambda_{t-1}}^t \log \frac{\hat{g}_N^{\lambda_{t-1}:i-1}(x_i)}{f_N(x_i)} = \sum_{i=\lambda_{t-1}+1}^t \log \frac{\hat{g}_N^{\lambda_{t-1}:i-1}(x_i)}{f_N(x_i)}.$$

Thus, if $k = \lambda_{t-1} + 1$ maximizes the sum $\sum_{i=k}^t \log \frac{\hat{g}_N^{\lambda_{t-1}:i-1}(x_i)}{f_N(x_i)}$, $k = \lambda_{t-1}$ also maximizes it. For this case, we choose the most likely change-point to be λ_{t-1} rather than $\lambda_{t-1} + 1$, as defined in (11). Note also that if $\lambda_{t-1} = t$, then we have $\tilde{S}_N(t) = \tilde{S}_N(t-1) = 0$, i.e., it takes more than one sample for our test statistic to move away from the zero boundary once it hits it. This is due to the way we define our estimator $\hat{g}_N^{\lambda_{t-1}:i-1}$ in (7), which at time t utilizes samples starting from the last most likely change-point $\lambda_{t-1} = t$ to the previous time $t-1$ to estimate the post-change distribution, i.e., it simply uses f_N as the estimator. From (11), we then have $\lambda_t = t$. Furthermore, if $\tilde{S}_N(t-1) = 0$ while $\tilde{S}_N(t) > 0$, then we must have $\lambda_{t-1} = t-1$.

The test statistic $\tilde{S}_N(t)$ can be efficiently computed using a recursive update as shown in the following result.

Theorem 1. For each $t \geq 0$, we have the update formula

$$\tilde{S}_N(t+1) = \max \left\{ \tilde{S}_N(t) + \log \frac{\hat{g}_N^{\lambda_t:t}(x_{t+1})}{f_N(x_{t+1})}, 0 \right\}, \quad (13)$$

$$\lambda_{t+1} = \begin{cases} \lambda_t & \text{if } \tilde{S}_N(t) + \log \frac{\hat{g}_N^{\lambda_t:t}(x_{t+1})}{f_N(x_{t+1})} > 0 \text{ or } \lambda_t = t+1, \\ t+2 & \text{otherwise,} \end{cases} \quad (14)$$

where $\tilde{S}_N(0) = 0$ and $\lambda_0 = 1$.

Proof: See Appendix A. ■

Similar to the CuSum test, the renewal property of the test statistic and the fact that it is non-negative implies that the worst case change-point ν for the ADD is at $\nu = 0$. We compare the performance of the stopping time (9) with that of (12) using simulations in Section V.

IV. PROPERTIES OF THE BG-CUSUM STATISTIC

In this section, we present some properties of the BG-CuSum statistic in order to give insights into the asymptotic behavior of the ARL and ADD of our test. The proofs for the results in this section are provided in Appendix B.

A. Estimating a lower bound for the ARL

In applications, a practitioner is required to set a threshold b for the problem of interest. Thus, it is of practical interest to have an estimate of the ARL of our test w.r.t. the threshold b . This is even more important in our context as it is not possible to set the threshold b w.r.t. the WADD since it varies with the unknown post-change distribution. In this subsection, we derive a lower bound for the ARL of the BG-CuSum test.

Define the stopping time

$$\zeta(b) = \inf\{t: \tilde{S}_N(t) \geq b \text{ or } (\tilde{S}_N(t) \leq 0 \text{ and } t > 2)\}.$$

Note that the condition $t > 2$ is used due to the lag in our test statistic; see the discussion after (12). We have the following lower bound for the ARL of the BG-CuSum test.

Proposition 1. For any threshold $b > 0$, we have

$$\text{ARL}(\tilde{\tau}(b)) = \frac{\mathbb{E}_\infty[\zeta(b)]}{\mathbb{P}_\infty(\tilde{S}_N(\zeta(b)) \geq b)} \geq e^b. \quad (15)$$

B. Growth rate of the BG-CuSum statistic

We study the error bounds of the growth rate of the BG-CuSum statistic $\tilde{S}_N$ under the assumption that the observed samples $X_1=x_1, X_2=x_2, \dots$ are generated by the post-change distribution g . We show that the growth rate $\tilde{S}_N(t)/t$ converges to the Kullback-Leibler divergence $D_{\text{KL}}(g_N \parallel f_N)$ r -quickly (see Section 2.4.3 of [29]) for $r=1$ under $\mathbb{P}_1$, which implies almost sure convergence. We then provide some heuristic insights into the asymptotic trade-off between the ARL and ADD of our BG-CuSum test as the threshold $b \rightarrow \infty$ in (12).

We first show that the regularized sample mean $\hat{g}_N^{1:i}$ in (7) is close to g_N with high probability when the sample size is large.

Proposition 2. *For any $\epsilon \in (0,1)$ and $x \in \mathbb{R}$, there exists $t_1 \in \mathbb{N}$ such that for all $i \geq t_1$, we have*

$$\mathbb{P}_1(|\hat{g}_N^{1:i}(x) - g_N(x)| \geq \epsilon) \leq 2e^{-i\epsilon^2/2}.$$

We next show that the instantaneous log-likelihood ratio $\log \frac{\hat{g}_N^{1:i}(x)}{f_N(x)}$ is close to the true log-likelihood ratio $\log \frac{g_N(x)}{f_N(x)}$, in the sense that the probability of a deviation of ϵ decreases to zero exponentially as the number of samples increases.

Proposition 3. *For any $\epsilon \in (0,1)$, there exists a $t_2 \in \mathbb{N}$ such that for all $i \geq t_2$, we have*

$$\mathbb{P}_1\left(\left|\log \frac{\hat{g}_N^{1:i}(X_{i+1})}{f_N(X_{i+1})} - \log \frac{g_N(X_{i+1})}{f_N(X_{i+1})}\right| \geq \epsilon\right) \leq 2e^{-i(g_N^{\min} \epsilon)^2/8}$$

where $g_N^{\min} = \min_x g_N(x)$.

Putting Propositions 2 and 3 together, we obtain the following theorem.

Proposition 4. *The empirical average $\frac{1}{t} \sum_{i=1}^t \log \frac{\hat{g}_N^{1:i-1}(X_i)}{f_N(X_i)}$ converges to $D_{\text{KL}}(g_N \parallel f_N)$ r -quickly for $r=1$ under the distribution $\mathbb{P}_1$ as $t \rightarrow \infty$.*

From Proposition 4, the probability of the growth rate of the BG-CuSum statistic $\tilde{S}_N(t)/t$ deviating from $D_{\text{KL}}(g_N \parallel f_N)$ by more than ϵ can be made arbitrarily small by increasing the number of samples t after the change point. Heuristically, this means that the ADD increases linearly at a rate of $1/D_{\text{KL}}(g_N \parallel f_N)$ w.r.t. the threshold b . Unfortunately, due to technical difficulties introduced by having to estimate the post-change distribution using $\hat{g}_N^{1:i-1}$, we are unable to quantify the asymptotic trade-off between ARL and ADD for our BG-CuSum test $\hat{\tau}$. In the following, we consider the asymptotic trade-off of the test $\hat{\tau}$ in (9) as an approximation of $\tilde{\tau}$, and use simulation results in Section V-C to verify that a similar trade-off applies for the BG-CuSum test.

Proposition 5. *The stopping time $\hat{\tau}(b)$ defined in (9) satisfies*

$$\text{ARL}(\hat{\tau}(b)) \geq e^b,$$

and

$$\text{WADD}(\hat{\tau}(b)) \leq \frac{b}{D_{\text{KL}}(g_N \parallel f_N)} + o(b) \quad \text{as } b \rightarrow \infty.$$

Furthermore if both f and g are discrete distributions (i.e., $p_0=q_0=0$), the stopping time $\hat{\tau}(b)$ is asymptotically optimal.

V. NUMERICAL RESULTS

In this section, we first compare the performance of our proposed BG-CuSum test with two other non-parametric change detection methods in the literature. We first perform simulations, and then we verify the performance of our method on real activity tracking data from [49].

A. Synthetic data

In our first set of simulations, we set the parameters $H=0$, $N=16$, $R=16$ and choose the pre-change distribution to be the standard normal distribution $\mathcal{N}(0,1)$. All the post-change distributions we use in our simulations are in $D(f,N)$ for this choice of parameters.

In our first experiment, we compare the performance of our method with [37], in which it is assumed that the change is an unknown shift in the location parameter. We also compare the performance of our method with [43], in which it is assumed that the means of the pre- and post-change distributions are different, and both distributions are unknown. In our simulation, we let the post-change distribution be $\mathcal{N}(\delta,1)$. We control the ARL at 500 and set the change-point $\nu=300$ for both methods while varying δ . The average detection delay is computed from 50,000 Monte Carlo trials and shown in Table I, where the smallest ADDs for each ARL are highlighted in boldface. We see that our method, despite not assuming that the change is a mean shift, achieves a comparable ADD as the method in [37]. Furthermore, we see that our method outperforms the method in [43].

TABLE I
ADD FOR POST-CHANGE DISTRIBUTION $\mathcal{N}(\delta,1)$ WITH ARL=500.

δ	0.125	0.75	1.5	2.25	3
Hawkins [37]	428.60	18.1	6.6	4.5	3.9
Darkhovskii [43]					
($N=280, b=0.25, c=0.1$)	458.71	26.15	13.02	8.55	6.58
BG-CuSum	344.78	17.9	6.6	3.2	2.3

We next consider a shift in variance for the post-change distribution. The method in [37], for example, will not be able to detect this change accurately as it assumes that the change in distribution is a shift in mean. Therefore, we also compare our method with the KS-CPM method [45], which is a non-parametric test that makes use of the Kolmogorov-Smirnov statistic to construct a sequential 2-sample test to test for a change-point. We control the ARL at 500 and change-point $\nu=300$ for all methods while varying δ for the post-change distribution $g \sim \mathcal{N}(0,\delta^2)$. The average detection delays computed from 50,000 Monte Carlo trials are shown in Table II. We see that our method outperforms both [37] and [45] in the ADD. We note that $D_{\text{KL}}(g_N \parallel f_N)$ decreases as δ varies from 0.2 to 1, and increases as δ varies from 1 to 2. Thus, the observation that the ADD increases as δ varies from 0.2 to 1, and decreases as δ varies from 1 to 2 agrees with our expectation.

Next, we test our method with the Laplace post-change distribution with the probability density function $g(x) = \frac{1}{2 \times 0.7071} e^{-\frac{|x|}{0.7071}}$. The location and scale parameter are chosen

TABLE II
ADD FOR POST-CHANGE DISTRIBUTION $\mathcal{N}(0, \delta^2)$ WITH ARL=500.

δ	0.2	0.33	0.5	1.5	2
Hawkins [37]	361.3	391.5	438.5	149.6	75.3
KS-CPM [45]	27.2	37.2	84.7	140.6	49.2
BG-CuSum	10.5	17.4	33.3	45.2	21.5

Fig. 1. Examples of test statistics $\hat{S}(t)$, $\tilde{S}(t)$ and $\hat{S}(t) - \tilde{S}(t)$ as a function of t for $D_{KL}(g_N \parallel f_N)=2$.

such that the first and second order moments of g and f are equal. We set the ARL=500 and two different values for the change-point ν . We performed 50,000 Monte Carlo trials to obtain Table III. The results show the method is able to identify the change from a normal to a Laplace distribution with the smallest ADD out of the three methods studied.

TABLE III
ADD FOR LAPLACE POST-CHANGE DISTRIBUTION WITH ARL=500.

ν	50	300
Hawkins [37]	592	828
KS-CPM [45]	284	217
BG-CuSum	156	154

In Fig. 1 and Fig. 2, we show how the BG-CuSum test statistic $\tilde{S}_N(t)$ behaves for two different post-change distributions. In Fig. 1, the pre-change distribution is $\mathcal{N}(0,1)$ and the post-change distribution is $\mathcal{N}(0.2,1)$. In Fig. 2, we consider distributions with both discrete and absolutely continuous components. The change is in the discrete component where the pre-change distribution is $0.5\mathcal{N}(0,1)+0.25\delta_{-1}+0.25\delta_1$ and the post-change distribution is $0.5\mathcal{N}(0,1)+0.33\delta_{-1}+0.17\delta_1$. We observe, in both cases, that $\tilde{S}_N(t)$ remains low during the pre-change regime and quickly rises in the post-change regime in both cases. Furthermore, the proposed recursive test statistics $\hat{S}_N$ is observed to track the test statistic $\tilde{S}_N$, which has known asymptotic properties as seen in Proposition 5.

B. Choice of N

Before applying the BG-CuSum test, the user has to choose an appropriate number of bins N . In the supplementary material [48], we have provided a procedure for

Fig. 2. Examples of test statistics $\hat{S}(t)$, $\tilde{S}(t)$ and $\hat{S}(t) - \tilde{S}(t)$ as a function of t when the change is in the discrete component such that the pre-change distribution is $0.5\mathcal{N}(0,1)+0.25\delta_{-1}+0.25\delta_1$ and post-change distribution is $0.5\mathcal{N}(0,1)+0.33\delta_{-1}+0.17\delta_1$.

Fig. 3. Graph of ADD against $\log(\text{ARL})$ for varying values of N with $f=\mathcal{N}(0,1)$ and $g=0.6\mathcal{N}(1,1)+0.4\mathcal{N}(-1,1)$.

determining a suitable choice of N . In this subsection, we present several results to illustrate the guiding principles for choosing N . First, we compare the ARL-ADD performance of BG-CuSum using different values of N . We performed the experiments with $\mathcal{N}(0,1)$ as the pre-change distribution, and Gaussian mixture model $0.6\mathcal{N}(1,1)+0.4\mathcal{N}(-1,1)$ as the post-change distribution for $N=2,4,8,16,32,64$, using 5000 Monte Carlo trials to obtain the ADD and ARL. The KL divergence $D_{KL}(g_N \parallel f_N)$ for respective values of N are 0.0094, 0.730, 0.1164, 0.1420, 0.1565 and 0.1645. From Fig. 3, we observe that the performance of BG-CuSum improves as N increases from 2 to 8 and degrades as N increases from 8 to 64. One reason for this is that the benefit of having a larger $D_{KL}(g_N \parallel f_N)$ does not out-weigh the larger number of samples required to accurately estimate the unknown post-change distribution g_N when N is large.

Fig. 4. Plot of $\Delta(b)$ against threshold b .

C. Asymptotic behaviour of the ARL and ADD of $\tilde{\tau}$

In Section IV, we showed that the growth of the BG-CuSum test statistic can be made arbitrarily close to $D_{KL}(g_N \| f_N)$ with a sufficiently large number of samples t . Heuristically, from Proposition 5, using $\hat{S}(t)$ in (8) as an approximation of the BG-CuSum test statistic $S_N(t)$, this implies that the average detection delay would grow at a rate of $1/D_{KL}(g_N \| f_N)$. To demonstrate this, we let $N=64$, $\mathcal{N}(0,1)$ as the pre-change distribution, and g to be one of several normal distributions with different means and variance 1 so that we have $D_{KL}(g_N \| f_N)=1, \frac{1}{2}, \frac{1}{3}, \frac{1}{4}, \frac{1}{5}$, respectively. Therefore we expect the asymptotic gradient of the ADD w.r.t. b to be 1, 2, 3, 4 and 5 respectively. We performed 1000 Monte Carlo trials to estimate the ADD for different values of b . Fig. 4 shows the plot of ADD against b and

$$\Delta(b) = \frac{\text{ADD}(\tau(b+h)) - \text{ADD}(\tau(b))}{h},$$

which approximates the gradient of WADD w.r.t. b . We used a step-size of $h=12$ to generate Fig. 4. We see that $\Delta(b)$ tends to 1, 2, 3, 4 and 5 respectively as b tends to infinity, which agrees with those predicted by our heuristic.

Next, we compare the asymptotic performance of $\tilde{\tau}$ in (12) and $\hat{\tau}$ in (9). We perform simulations with $N=4$ and $f=\mathcal{N}(0,1)$ and g to be one of several normal distributions with different means and variance 1 so that we have $D_{KL}(g_N \| f_N)=1, 2$ and 3. We perform 5000 Monte Carlo trials to estimate the ARL and ADD of both $\hat{\tau}$ and $\tilde{\tau}$. Fig. 5 shows that the ADD of $\tilde{\tau}$ remains close to $\hat{\tau}$ as the ARL becomes large.

D. WISDM Actitracker Dataset

We now apply our BG-CuSum test on real activity tracking data using the WISDM Actitracker dataset [49], which contains 1,098,207 samples from the accelerometers of Android phones. The dataset was collected from 29 volunteer subjects who each carried an Android phone in their front leg pocket while performing a set of activities such as walking, jogging, ascending stairs, descending stairs, sitting and standing. The

Fig. 5. Comparison of the trade-off between ADD and $\log(\text{ARL})$ for $\tilde{\tau}$ and $\hat{\tau}$.

signals from the accelerometers were collected at 20 samples per second. The movements generated when each activity by each individual can be assumed to be consistent across time. Hence we assume that the samples generated are i.i.d. for each activity.

To test the effectiveness of the BG-CuSum test on real data, we apply it to detect the changes in the subject's activity. We aim to detect the change from walking activity to any other activity. There are 45 segments in the dataset in which there is a switch from walking to other activities. For each of these segments, we use the first half of the samples from the walking activity period to learn the pre-change pdf f for the data.

As each sample at time t from the phone's accelerometer is a 3-dimensional vector $\mathbf{a}(t)$, we perform change detection on the sequence $x_t = \|\mathbf{a}(t)\|_2$ instead. We set the number of bins in the BG-CuSum test to be $H=0$ and $N=32$. Using the first T samples from each walking activity segment, we estimate the boundaries of the intervals I_j^N , such that $\int_{I_j^N} f(x) dx = 1/N$ by setting $I_1^N = (-\infty, x_{(\lfloor T/N \rfloor)}]$, $I_j^N = (x_{(\lfloor (j-1)T/N \rfloor)}, x_{(\lfloor jT/N \rfloor)}]$ for $1 < j < N-1$, and $I_N^N = (x_{(\lfloor (N-1)T/N \rfloor)}, \infty)$, where $x_{(n)}$ is the n -th order statistic of x_t .

In order to control the ARL of the BG-CuSum test to be 6000, we set the threshold b to be 2.69. The ADD for BG-CuSum on the WISDM Actitracker dataset is found to be about 30.2.

In Fig. 6, we present some examples of the performance of BG-CuSum. Similar to the simulations, we observe that in both cases, the test statistic $\hat{S}_N(t)$ remains low during the pre-change regime and quickly rises in the post-change regime. In Fig. 7, we present the notched boxplots of the detection delays of different algorithms when the ARL is set at 6000. We observe that the BG-CuSum test out-performs both the KS-CPM [45] and Hawkins [37] algorithm.

VI. CONCLUSION

We have studied the sequential change detection problem when the pre-change distribution f is known, and the post-

Fig. 6. Examples of trial performed for pre-change activity of walking, and post-change activity of (a) jogging and (b) ascending upstairs. The black dotted line indicates the boundary between the pre-change and post-change regimes.

Fig. 7. Box-plot of the detection delays of each algorithm when the ARL is 6000.

change distribution g lies in a family of distributions with k -th moment differing from f by at least ϵ . We proposed a sequential change detection method that partitions the sample

space into bins, and a test statistic that can be updated recursively in real time with low complexity. We analyzed the growth rate of our test statistic and used it to heuristically deduce the asymptotic relationship between the ADD and the ARL. Tests on both synthetic and real data suggest that our proposed BG-CuSum test outperforms several other non-parametric tests in the literature. Furthermore, simulations indicate that the BG-CuSum test approaches the performance of $\hat{\tau}$, which has known asymptotic properties, as the ARL becomes large. We provided a lower bound on the ARL of the BG-CuSum test to aid the setting of the threshold b . One direction for future work would be to derive the WADD for the BG-CuSum test. This remains an open research problem due to the technical difficulties introduced by having to estimate the post-change distribution using an estimated change-point.

Although we have assumed that the pre- and post-change distributions are defined on $\mathbb{R}$, the BG-CuSum test derived in this paper can be applied to cases where f and g are generalized pdfs on $\mathbb{R}^n$ with $n > 1$. To see how this can be done, we assume that $g \in D(f, N)$, where $N = N_0^n$ is a power of n . We can then divide $\mathbb{R}^n$ into N equi-probable sets w.r.t. f_c by sequentially dividing each dimension into N_0 equi-probable intervals. After obtaining these N sets, we can apply the BG-CuSum test directly. The results derived in Section IV can also be extended to the case where the distributions are on $\mathbb{R}^n$. However, the amount of data required to learn the pre-change distribution f increases quickly w.r.t. n , which limits its application in practice.

The assumption that the post change distribution g is absolutely continuous w.r.t. f can be further relaxed. For example, the results developed in Section III and IV extend to the case with a continuous pre-change distribution f and a discrete post-change distribution g . The recursive update scheme in Theorem 1 and the lower bound on the ARL in Proposition 1 still holds. However, as g is not absolutely continuous w.r.t. f , the instantaneous empirical log-likelihood ratio $\log \frac{\hat{g}_N^{1:i-1}(X_i)}{f_N(X_i)} = \infty$ and the instantaneous log-likelihood ratio $\log \frac{\hat{g}_N(X_i)}{f_N(X_i)} = \infty$ with positive probability under $\mathbb{P}_1$. Thus, we have that $D_{\text{KL}}(g_N \parallel f_N) = \infty$. Furthermore, in Proposition 4, the empirical average $\frac{1}{t} \sum_{i=1}^t \log \frac{\hat{g}_N^{1:i-1}(X_i)}{f_N(X_i)}$ diverges to infinity almost surely.

One possible direction for future research is to extend the approach to quickest change detection in Hidden Markov Models [50], [51] when the post-change transition probability matrix is unknown. In order to compute the likelihood function, a similar estimation scheme for the transition probability matrix can be derived using the maximum-likelihood state estimates of the observed samples. However, more work is required to derive the asymptotic operating characteristics of the stopping time.

APPENDIX A PROOF OF THEOREM 1

In order to show that $\tilde{S}_N$ can be computed recursively, we require the following lemmas.

Lemma A.1. Suppose $\lambda_p = \lambda_{p+1} = p+1$ and $\tilde{S}_N(p+2), \dots, \tilde{S}_N(p+n) > 0$ for some $n \geq 1$, then we have

$$\lambda_{p+2} = \lambda_{p+3} = \dots = \lambda_{p+n} = p+1.$$

Proof: We prove by contradiction. Suppose that there exists $t \in \{2, \dots, n\}$ such that $\lambda_{p+t} \neq p+1$. Let t_0 be the smallest of all such indices t . Since $\lambda_{p+t_0} \neq p+1$, following the definition of λ_{p+t_0} , we obtain $\lambda_{p+t_0} > \lambda_{p+t_0-1} = p+1$. Furthermore, by (11), $\lambda_{p+t_0} \neq \lambda_{p+t_0-1} + 1 = p+2$. Thus we have $\lambda_{p+t_0} > p+2$. From (10), we have

$$\begin{aligned} \tilde{S}_N(p+t_0) &= \max_{\lambda_{p+t_0-1} \leq k \leq p+t_0+1} \sum_{i=k}^{p+t_0} \log \frac{\hat{g}_N^{\lambda_{p+t_0-1}:i-1}(x_i)}{f_N(x_i)} \\ &= \max_{p+1 \leq k \leq p+t_0+1} \sum_{i=k}^{p+t_0} \log \frac{\hat{g}_N^{p+1:i-1}(x_i)}{f_N(x_i)} \\ &= \sum_{i=\lambda_{p+t_0}}^{p+t_0} \log \frac{\hat{g}_N^{p+1:i-1}(x_i)}{f_N(x_i)}, \end{aligned} \quad (16)$$

where the last equality follows from the definition (11). We then obtain

$$\sum_{i=p+1}^{\lambda_{p+t_0}-1} \log \frac{\hat{g}_N^{p+1:i-1}(x_i)}{f_N(x_i)} = \sum_{i=p+1}^{p+t_0} \log \frac{\hat{g}_N^{p+1:i-1}(x_i)}{f_N(x_i)} - \tilde{S}_N(p+t_0) \leq 0, \quad (17)$$

where the inequality follows from (16). Since $p+1 < \lambda_{p+t_0} - 1 \leq p+t_0$, let $\lambda_{p+t_0}-1 = p+t_1$ where $2 \leq t_1 \leq t_0 \leq n$. Since t_0 is the smallest t such that $\lambda_{p+t} \neq p+1$, we have $\lambda_{p+t_1-1} = p+1$ and

$$\begin{aligned} \tilde{S}_N(p+t_1) &= \sum_{i=\lambda_{p+t_1}}^{p+t_1} \log \frac{\hat{g}_N^{\lambda_{p+t_1}-1:i-1}(x_i)}{f_N(x_i)} \\ &= \sum_{i=\lambda_{p+t_0}}^{\lambda_{p+t_0}-1} \log \frac{\hat{g}_N^{p+1:i-1}(x_i)}{f_N(x_i)} \leq 0, \end{aligned} \quad (18)$$

where the last inequality follows trivially if $t_1 = t_0$ and from (17) if $t_1 < t_0$. This inequality contradicts our assumption that $\tilde{S}_N(p+2), \dots, \tilde{S}_N(p+n) > 0$. The lemma is now proved. ■

Lemma A.2. If $\tilde{S}_N(t+1) > 0$, then $\tilde{S}_N(t+1) = \tilde{S}_N(t) + \log \frac{\hat{g}_N^{\lambda_t:t}(x_{t+1})}{f_N(x_{t+1})}$

Proof: We first consider the case where $\tilde{S}_N(t) = 0$. Since $\tilde{S}_N(t+1) > 0$, we must have $\lambda_t = t$. We then obtain

$$\begin{aligned} \tilde{S}_N(t+1) &= \max_{\lambda_t \leq k \leq t+2} \sum_{i=k}^{t+1} \log \frac{\hat{g}_N^{\lambda_t:i-1}(x_i)}{f_N(x_i)} \\ &= \max_{t \leq k \leq t+2} \sum_{i=k}^{t+1} \log \frac{\hat{g}_N^{t:i-1}(x_i)}{f_N(x_i)} \\ &= \log \frac{\hat{g}_N^{t:t}(x_{t+1})}{f_N(x_{t+1})} \\ &= \tilde{S}_N(t) + \log \frac{\hat{g}_N^{\lambda_t:t}(x_{t+1})}{f_N(x_{t+1})}. \end{aligned}$$

For the case where $\tilde{S}_N(t) > 0$, let $n \geq 0$ be the largest integer such that $\tilde{S}_N(t-n), \dots, \tilde{S}_N(t) > 0$. If $n=0$, we trivially

have $\lambda_{t-1} = \lambda_t = t-1$. If $n > 0$, since $\tilde{S}_N(t-n-1) = 0$, we have $\lambda_{t-n-1} = t-n-1$ and Lemma A.1 yields $\lambda_{t-n} = \dots = \lambda_t = t-n-1$. We then obtain

$$\begin{aligned} \tilde{S}_N(t+1) &= \max_{\lambda_t \leq k \leq t+1} \sum_{i=k}^{t+1} \log \frac{\hat{g}_N^{\lambda_t:i-1}(x_i)}{f_N(x_i)} \\ &= \max_{\lambda_t \leq k \leq t+1} \left\{ \sum_{i=k}^t \log \frac{\hat{g}_N^{\lambda_t:i-1}(x_i)}{f_N(x_i)} + \log \frac{\hat{g}_N^{\lambda_t:t}(x_{t+1})}{f_N(x_{t+1})} \right\} \\ &= \max_{\lambda_{t-1} \leq k \leq t+1} \left\{ \sum_{i=k}^t \log \frac{\hat{g}_N^{\lambda_{t-1}:i-1}(x_i)}{f_N(x_i)} \right\} + \log \frac{\hat{g}_N^{\lambda_t:t}(x_{t+1})}{f_N(x_{t+1})} \\ &= \tilde{S}_N(t) + \log \frac{\hat{g}_N^{\lambda_t:t}(x_{t+1})}{f_N(x_{t+1})}, \end{aligned}$$

and the proof is complete. ■

We are now ready to show that the statistic $\tilde{S}_N$ can be updated recursively. From Lemma A.2 and noting that $\tilde{S}_N(t) \geq 0$ for all t , (13) follows immediately. If $\tilde{S}_N(t+1) = 0$ and $\lambda_t \neq t+1$, then $\lambda_{t+1} = t+2$ from (11). Next, if $\tilde{S}_N(t+1) = 0$ and $\lambda_t = t+1$, then $\lambda_{t+1} = t+1 = \lambda_t$ from (11). Finally, if $\tilde{S}_N(t+1) > 0$, then by Lemma A.1, we have $\lambda_{t+1} = \lambda_t$, and (14) follows. The proof is now complete.

APPENDIX B

PROOFS OF RESULTS IN SECTION IV

In this appendix, we let $M = N + H$ and Y_k^j for $k=1, 2, \dots$ to be i.i.d. random variables such that

$$Y_k^j = \begin{cases} 1 & \text{if } X_k \in I_j^N, \\ 0 & \text{if } X_k \notin I_j^N. \end{cases}$$

Recalling (6), we then have

$$g_N(x) = \mathbb{E} \left[Y_1^j \right]$$

and the regularized sample mean $\hat{g}_N^{1:i}(x)$ defined in (7) can be written as

$$\hat{g}_N^{1:i}(x) = \frac{\sum_{k=1}^i Y_k^j + R}{i + MR}$$

where j is the unique integer such that $x \in I_j^N$. We first study the error bounds for the regularized sample mean $\hat{g}_N^{1:i}(x)$. We then derive error bounds for the instantaneous sample log-likelihood ratio $\log \frac{\hat{g}_N^{1:i}(x)}{f_N(x)}$. Finally we combine all the results together to derive an error bound for the growth rate of the BG-CuSum statistic.

A. Proof of Proposition 1

Following arguments identical to those in Chapter 2 of [52], it can be shown that

$$\text{ARL}(\tilde{\tau}(b)) = \frac{\mathbb{E}_\infty[\zeta(b)]}{\mathbb{P}_\infty(\tilde{S}_N(\zeta(b)) \geq b)}. \quad (19)$$

In order to obtain a lower bound for the ARL, we first note that for any $b > 0$, we have

$$\mathbb{E}_\infty[\zeta(b)] \geq \mathbb{E}_\infty[\zeta(0)] = 1. \quad (20)$$

We also have

$$\mathbb{P}_\infty(\tilde{S}_N(\zeta(b)) \geq b) = \mathbb{P}_\infty\left(\sum_{i=1}^{\zeta(b)} \log \frac{\hat{g}_N^{1:i-1}(X_i)}{f_N(X_i)} \geq b\right) \quad (21)$$

$$\begin{aligned} &= \mathbb{P}_\infty\left(\prod_{i=1}^{\zeta(b)} \frac{\hat{g}_N^{1:i-1}(X_i)}{f_N(X_i)} \geq e^b\right) \\ &= \sum_{t=1}^{\infty} \mathbb{P}_\infty(E_t), \end{aligned} \quad (22)$$

where E_t is the event $\left\{\prod_{i=1}^t \frac{\hat{g}_N^{1:i-1}(X_i)}{f_N(X_i)} \geq e^b, \zeta(b)=t\right\}$. The equality in (21) follows from Theorem 1 and noting that since $\zeta(b)$ is the first time $\tilde{S}_N(t)$ exceeds b or falls below 0, $\lambda_i = \lambda_0 = 1$ for all $i=1, \dots, \zeta(b)$ under the event $\tilde{S}_N(\zeta(b)) \geq b$.

Let $J_i \in \{1, \dots, N+H\}$ denote the index of the bin in $\{I_j^N\}_{j=1}^{N+H}$ that X_i falls into. Let F_t be the event such that $\{J_i\}_{i=1}^t \in F_t$ if and only if $\{X_i\}_{i=1}^t \in E_t$. For any $i \geq 1$ and any sequence $(x_1, \dots, x_i)$ with corresponding bin indices $(j_1, \dots, j_i)$, let $q(j_i | j_{1:i-1}) = \hat{g}_N^{1:i-1}(x_i)$. From the Kolmogorov Extension Theorem [53], there exists a probability measure $\mathbb{Q}$ on $\{1, \dots, N+H\}^{\mathbb{N}}$ such that $\mathbb{Q}(J_1=j_1, \dots, J_t=j_t) = \prod_{i=1}^t q(j_i | j_{1:i-1})$ for all $t \geq 1$. We then have

$$\begin{aligned} \mathbb{P}_\infty(E_t) &= \mathbb{E}_\infty[\mathbb{1}_{E_t}] \\ &\leq \mathbb{E}_\infty\left[e^{-b} \prod_{i=1}^t \frac{\hat{g}_N^{1:i-1}(X_i)}{f_N(X_i)} \mathbb{1}_{E_t}\right] \\ &\leq e^{-b} \sum_{\{j_i\} \in F_t} \prod_{i=1}^t q(j_i | j_{1:i-1}) \\ &= e^{-b} \mathbb{Q}(F_t). \end{aligned}$$

From (22), we obtain

$$\begin{aligned} \mathbb{P}_\infty(\tilde{S}_N(\zeta(b)) \geq b) &\leq e^{-b} \sum_{t=1}^{\infty} \mathbb{Q}(F_t) \\ &\leq e^{-b}. \end{aligned} \quad (23)$$

where the final inequality follows because $\{F_t\}_{t \in \mathbb{N}}$ are mutually exclusive events. Thus, from (19), (20) and (23), we have

$$\text{ARL}(\tilde{\tau}(b)) = \frac{\mathbb{E}_\infty[\zeta(b)]}{\mathbb{P}_\infty(\tilde{S}_N(\zeta(b)) \geq b)} \geq e^b. \quad (24)$$

The proof is now complete.

B. Proof of Proposition 2

In the design of the BG-CuSum test statistic, we replace the maximum likelihood estimator $g_N^{1:i}$ with its regularized sample mean $\hat{g}_N^{1:i}$. In this subsection, we derive a bound for the probability that $\hat{g}_N^{1:i}(x)$ deviates from $g_N(x)$ by at least $\epsilon > 0$. We start with a few elementary lemmas.

Lemma B.1. For any $\epsilon \geq 0$, $N \geq 1$, and $x \in \mathbb{R}$, we have

$$\mathbb{P}_1\left(\left|\hat{g}_N^{1:i}(x) - \frac{\mathbb{E}[\sum_{k=1}^i Y_k^j] + R}{i+MR}\right| \geq \epsilon\right) \leq 2e^{-2i\epsilon^2}$$

where j is the unique integer such that $x \in I_j^N$.

Proof: By applying Hoeffding's inequality [54], we obtain

$$\begin{aligned} &\mathbb{P}_1\left(\left|\hat{g}_N^{1:i}(x) - \frac{\mathbb{E}[\sum_{k=1}^i Y_k^j] + R}{i+MR}\right| \geq \epsilon\right) \\ &= \mathbb{P}_1\left(\left|\frac{\sum_{k=1}^i Y_k^j + R}{i+MR} - \frac{\mathbb{E}[\sum_{k=1}^i Y_k^j] + R}{i+MR}\right| \geq \epsilon\right) \\ &= \mathbb{P}_1\left(\left|\frac{\sum_{k=1}^i Y_k^j + R}{i} - \frac{\mathbb{E}[\sum_{k=1}^i Y_k^j] + R}{i}\right| \geq \frac{i+MR}{i}\epsilon\right) \\ &\leq 2e^{-2i\left(\frac{i+MR}{i}\right)^2 \epsilon^2} \leq 2e^{-2i\epsilon^2}, \end{aligned}$$

and the proof is complete. $\blacksquare$

Lemma B.2. For any $\epsilon \in (0, 1)$ and $x \in \mathbb{R}$, there exists $i_0 \in \mathbb{N}$ such that for all $i \geq i_0$, we have

$$\left|\frac{\mathbb{E}[\sum_{k=1}^i Y_k^j] + R}{i+MR} - g_N(x)\right| < \epsilon$$

where j is the unique integer such that $x \in I_j^N$.

Proof: For $i \geq (1-\epsilon)MR/\epsilon$, we have

$$\begin{aligned} \left|\frac{\mathbb{E}[\sum_{k=1}^i Y_k^j] + R}{i+MR} - g_N(x)\right| &= \left|\frac{ig_N(x) + R}{i+MR} - g_N(x)\right| \\ &= \left|\frac{R - MRg_N(x)}{i+MR}\right| \\ &= \frac{R}{i+MR} |1 - Mg_N(x)| \\ &\leq \frac{MR}{i+MR} \\ &\leq \epsilon, \end{aligned}$$

and the proof is complete. $\blacksquare$

Putting everything together, we now proceed to the proof of Proposition 2. Given $0 < \epsilon < 1$, by Lemma B.2 there exists $i_1 > 0$ such that for all $i \geq i_1$, we have

$$\left|\frac{\mathbb{E}[\sum_{k=1}^i Y_k^j] + R}{i+MR} - g_N(x)\right| < \frac{\epsilon}{2}.$$

Therefore, we obtain

$$\begin{aligned} &\mathbb{P}_1(|\hat{g}_N^{1:i}(x) - g_N(x)| \geq \epsilon) \\ &= \mathbb{P}_1\left(\left|\hat{g}_N^{1:i}(x) - \frac{\mathbb{E}[\sum_{k=1}^i Y_k^j] + R}{i+MR} + \frac{\mathbb{E}[\sum_{k=1}^i Y_k^j] + R}{i+MR} - g_N(x)\right| \geq \epsilon\right) \\ &\leq \mathbb{P}_1\left(\left|\hat{g}_N^{1:t}(x) - \frac{\mathbb{E}[\sum_{k=1}^i Y_k^j] + R}{i+MR}\right| + \left|\frac{\mathbb{E}[\sum_{k=1}^i Y_k^j] + R}{i+MR} - g_N(x)\right| \geq \epsilon\right) \\ &\leq \mathbb{P}_1\left(\left|\hat{g}_N^{1:i}(x) - \frac{\mathbb{E}[\sum_{k=1}^i Y_k^j] + R}{i+MR}\right| \geq \frac{\epsilon}{2}\right) \\ &\leq 2e^{-\frac{1}{2}i\epsilon^2}, \end{aligned}$$

where the last inequality follows from Lemma B.1, and the proposition is proved.

C. Proof of Proposition 3

In this subsection, we use previous results on the regularized sample mean $\hat{g}_N^{1:i}$ to study the instantaneous log-likelihood ratio $\log \frac{\hat{g}_N^{1:i}(x)}{f_N(x)}$ used in the BG-CuSum test statistic.

Lemma B.3. For any $\epsilon \in (0,1)$, there exists a $t_2 \in \mathbb{N}$ such that for all $i \geq t_2$ and any $x \in \mathbb{R}$, we have

$$\mathbb{P}_1 \left(\left| \log \hat{g}_N^{1:i}(x) - \log g_N(x) \right| \geq \epsilon \right) \leq 2e^{-i(g_N^{\min} \epsilon)^2 / 8}$$

where $g_N^{\min} = \min_x g_N(x)$.

Proof:

Using the inequality $\log x \leq x-1$, we obtain for each $x \in \mathbb{R}$,

$$\mathbb{P}_1 \left(\log \frac{\hat{g}_N^{1:i}(x)}{g_N(x)} \geq \epsilon \right) \leq \mathbb{P} \left(\hat{g}_N^{1:i}(x) - g_N(x) \geq \epsilon g_N(x) \right)$$

and

$$\mathbb{P}_1 \left(\log \frac{g_N(x)}{\hat{g}_N^{1:i}(x)} \geq \epsilon \right) \leq \mathbb{P} \left(g_N(x) - \hat{g}_N^{1:i}(x) \geq \frac{\epsilon}{1+\epsilon} g_N(x) \right).$$

Therefore, we have

$$\begin{aligned} & \mathbb{P}_1 \left(\left| \log \hat{g}_N^{1:i}(x) - \log g_N(x) \right| \geq \epsilon \right) \\ &= \mathbb{P}_1 \left(\log \frac{\hat{g}_N^{1:i}(x)}{g_N(x)} \geq \epsilon \right) + \mathbb{P}_1 \left(\log \frac{g_N(x)}{\hat{g}_N^{1:i}(x)} \geq \epsilon \right) \\ &\leq \mathbb{P}_1 \left(\hat{g}_N^{1:i}(x) - g_N(x) \geq \epsilon g_N(x) \right) \\ &\quad + \mathbb{P}_1 \left(g_N(x) - \hat{g}_N^{1:i}(x) \geq \frac{\epsilon}{1+\epsilon} g_N(x) \right) \\ &\leq \mathbb{P}_1 \left(\left| \hat{g}_N^{1:i}(x) - g_N(x) \right| \geq \min \left(\epsilon g_N(x), \frac{\epsilon}{1+\epsilon} g_N(x) \right) \right) \\ &\leq \mathbb{P}_1 \left(\left| \hat{g}_N^{1:i}(x) - g_N(x) \right| \geq \frac{1}{2} \epsilon g_N(x) \right) \\ &\leq 2e^{-i(g_N(x)\epsilon)^2 / 8} \leq 2e^{-i(g_N^{\min} \epsilon)^2 / 8}, \end{aligned}$$

where the last inequality follows from Proposition 2 for all $i \geq t_1^x$, where $t_1^j \in \mathbb{N}$ is chosen to be sufficiently large. Taking $t_2 = \max_{x \in \mathbb{R}} t_1^x$, the lemma follows. ■

We now proceed to the proof of Proposition 3. Using the law of total probability and Lemma B.3, we have for $i \geq t_2$, where t_2 is as given in Lemma B.3,

$$\begin{aligned} & \mathbb{P}_1 \left(\left| \log \frac{\hat{g}_N^{1:i}(X_{i+1})}{f_N(X_{i+1})} - \log \frac{g_N(X_{i+1})}{f_N(X_{i+1})} \right| \geq \epsilon \right) \\ &= \sum_{j=1}^M \mathbb{P}_1 \left(\left| \log \hat{g}_N^{1:i}(X_{i+1}) - \log g_N(X_{i+1}) \right| \geq \epsilon \mid X_{i+1} \in I_j^N \right) \\ &\quad \times \mathbb{P}_1 \left(X_{i+1} \in I_j^N \right) \\ &= \sum_{j=1}^M \mathbb{P}_1 \left(\left| \log \hat{g}_N^{1:i}(X_{i+1}) - \log g_N(X_{i+1}) \right| \geq \epsilon \mid X_{i+1} \in I_j^N \right) \\ &\quad \times \mathbb{P}_1 \left(X_{i+1} \in I_j^N \right) \\ &= \sum_{j=1}^M \mathbb{P}_1 \left(\left| \log \hat{g}_N^{1:i}(X_{i+1}) - \log g_N(X_{i+1}) \right| \geq \epsilon \right) \\ &\quad \times \mathbb{P}_1 \left(X_{i+1} \in I_j^N \right) \\ &\leq \sum_{j=1}^M 2e^{-\frac{i}{8}(g_N^{\min} \epsilon)^2} \mathbb{P}_1 \left(X_{i+1} \in I_j^N \right) \\ &= 2e^{-i(g_N^{\min} \epsilon)^2 / 8}, \end{aligned}$$

and the proof is complete.

D. Proof of Proposition 4

In this subsection, we use results derived in the previous two subsections to study the growth rate $\tilde{S}_N(t)/t$ of the BG-CuSum test statistic.

Lemma B.4. For any $\epsilon \in (0,1)$, there exists a $t_3 \in \mathbb{N}$ such that for all $t \geq t_3$, we have

$$\mathbb{P}_1 \left(\left| \frac{1}{t} \sum_{i=1}^t \log \frac{\hat{g}_N^{1:i-1}(X_i)}{g_N(X_i)} \right| \geq \epsilon \right) \leq c_2 e^{-c_1 t^{1-\epsilon}},$$

where c_1, c_2 are positive constants.

Proof: Let $l = \lceil t^{1-\epsilon} \rceil$. For all $i \leq l$, we have from (7),

$$\hat{g}_N^{1:i-1}(X_i) \geq \frac{R}{l+MR} \quad \mathbb{P}_1\text{-a.s.},$$

which yields

$$\left| \log \hat{g}_N^{1:i-1}(X_i) - \log g_N(X_i) \right| \leq \log \frac{l+MR}{R} + |\log g_N^{\min}|, \quad (25)$$

where $g_N^{\min} = \min_j g_N(j) > 0$ since $f_N(j) = 1/N$ for all j and is absolutely continuous w.r.t. g_N . There exists $t_3 \in \mathbb{N}$ such that for all $t \geq t_3$,

$$\begin{aligned} & \frac{l}{t} \left(\log \frac{l+MR}{R} + |\log g_N^{\min}| \right) \\ &\leq \frac{t^{1-\epsilon} + 1}{t} \left(\log \frac{t^{1-\epsilon} + 1 + MR}{R} + |\log g_N^{\min}| \right) \\ &\leq \left(t^{-\epsilon} + \frac{1}{t} \right) \left(\log \frac{t^{1-\epsilon} + 1 + MR}{R} + |\log g_N^{\min}| \right) \leq \frac{\epsilon}{2}. \end{aligned}$$

We then obtain for $t \geq t_3$,

$$\begin{aligned} & \mathbb{P}_1 \left(\left| \frac{1}{t} \sum_{i=1}^t \log \frac{\hat{g}_N^{1:i-1}(X_i)}{g_N(X_i)} \right| \geq \epsilon \right) \\ &\leq \mathbb{P}_1 \left(\left| \frac{1}{t-l} \sum_{i=l+1}^t \log \frac{\hat{g}_N^{1:i-1}(X_i)}{g_N(X_i)} \right| + \frac{l}{t} \left(\log \frac{l+MR}{R} + |\log g_N^{\min}| \right) \geq \epsilon \right) \\ &\leq \mathbb{P}_1 \left(\left| \frac{1}{t-l} \sum_{i=l+1}^t \log \frac{\hat{g}_N^{1:i-1}(X_i)}{g_N(X_i)} \right| \geq \frac{\epsilon}{2} \right) \\ &\leq \sum_{i=l+1}^t \mathbb{P}_1 \left(\left| \log \hat{g}_N^{1:i-1}(X_i) - \log g_N(X_i) \right| \geq \frac{\epsilon}{2} \right) \\ &\leq 2 \sum_{i=l+1}^t e^{-i(g_N^{\min} \epsilon)^2 / 8} \\ &\leq c_2 e^{-c_1 l} \leq c_2 e^{-c_1 t^{1-\epsilon}}, \end{aligned}$$

where the penultimate inequality follows from Proposition 3, $c_1 = (g_N^{\min} \epsilon)^2 / 8$, and $c_2 = 2(1 - e^{-c_1})^{-1}$. The lemma is now proved. ■

Finally, we are ready to prove Proposition 4. We begin by showing that $\frac{1}{t} \sum_{i=1}^t \log \frac{\hat{g}_N^{1:i-1}(X_i)}{g_N(X_i)}$ converges to zero r-quickly under the distribution $\mathbb{P}_1$. For any $\epsilon \in (0,1)$, let

$$L_\epsilon = \sup \left\{ t : \left| \frac{1}{t} \sum_{i=1}^t \log \frac{\hat{g}_N^{1:i-1}(X_i)}{g_N(X_i)} \right| > \epsilon \right\}.$$

We have

$$\begin{aligned}\mathbb{E}_1[L_\epsilon] &= \sum_{n=1}^{\infty} \mathbb{P}_1(L_\epsilon \geq n) \\ &= \sum_{n=1}^{\infty} \mathbb{P}_1\left(\left|\frac{1}{t} \sum_{i=1}^t \log \frac{\hat{g}_N^{1:i-1}(X_i)}{g_N(X_i)}\right| \geq \epsilon, \text{ for some } t \geq n\right) \\ &\leq \sum_{n=1}^{\infty} \sum_{t=n}^{\infty} \mathbb{P}_1\left(\left|\frac{1}{t} \sum_{i=1}^t \log \frac{\hat{g}_N^{1:i-1}(X_i)}{g_N(X_i)}\right| \geq \epsilon\right) \\ &= \sum_{n=1}^{\infty} \sum_{t=n}^{\infty} c_2 e^{-c_1 t^{1-\epsilon}} \\ &= c_2 \sum_{n=1}^{\infty} n e^{-c_1 n^{1-\epsilon}} < \infty,\end{aligned}$$

where the penultimate equality follows from Lemma B.4, and c_1, c_2 are positive constants. Thus, $\frac{1}{t} \sum_{i=1}^t \log \frac{\hat{g}_N^{1:i-1}(X_i)}{g_N(X_i)}$ converges to zero r -quickly for $r=1$ under the distribution $\mathbb{P}_1$.

Since $\mathbb{E}_1\left[\left(\log \frac{g_N(X)}{f_N(X)}\right)^2\right] < \infty$, from Theorem 2.4.4 of [29], $\frac{1}{t} \sum_{i=1}^t \log \frac{g_N(X_i)}{f_N(X_i)}$ converges to $D_{\text{KL}}(g_N \parallel f_N)$ r -quickly for $r=1$ under the distribution $\mathbb{P}_1$. Therefore, $\frac{1}{t} \sum_{i=1}^t \log \frac{\hat{g}_N^{1:i-1}(X_i)}{f_N(X_i)}$ converges to $D_{\text{KL}}(g_N \parallel f_N)$ r -quickly for $r=1$ under $\mathbb{P}_1$ and the proof is complete.

E. Proof of Proposition 5

Let

$$\hat{T} = \inf\left\{t: \sum_{i=1}^t \log \frac{\hat{g}_N^{1:i-1}(X_i)}{f_N(X_i)} > b\right\}.$$

Using Proposition 4 and Corollary 3.4.1 in [29], we obtain

$$\mathbb{E}_1[\hat{T}] \sim \frac{b}{D_{\text{KL}}(g_N \parallel f_N)} \quad \text{as } b \rightarrow \infty.$$

Using arguments similar to those that led to Eq (23), we obtain

$$\mathbb{P}_\infty(\hat{T} < \infty) = \sum_{t=1}^{\infty} \mathbb{P}_\infty(\hat{T} = t) \leq e^{-b}.$$

Applying results from Theorem 6.16 in [28] to translate our understanding of $\hat{T}$ onto $\hat{\tau}(b)$, we obtain

$$\text{ARL}(\hat{\tau}(b)) \geq \frac{1}{\mathbb{P}_\infty(\hat{T} < \infty)} = e^b$$

and that

$$\text{WADD}(\hat{\tau}(b)) \leq \mathbb{E}_1[\hat{T}] \sim \frac{b}{D_{\text{KL}}(g_N \parallel f_N)} \quad \text{as } b \rightarrow \infty.$$

For the case where f and g are discrete distributions, we have $f_N = f$ and $g_N = g$, so that these bounds coincide with the bounds for the CuSum stopping time when both f and g are known. Thus, for this case, the test is asymptotically optimal and the theorem is now proved.

REFERENCES

- [1] Y. C. Chen, T. Banerjee, A. D. Domínguez-García, and V. V. Veeravalli, "Quickest line outage detection and identification," *IEEE Trans. on Power Syst.*, vol. 31, no. 1, pp. 749–758, 2016.
- [2] G. Rovatsos, X. Jiang, A. D. Domínguez-García, and V. V. Veeravalli, "Statistical power system line outage detection under transient dynamics," *IEEE Trans. on Signal. Process.*, vol. 65, no. 11, pp. 2787–2797, 2017.
- [3] V. M. R. Muggeo and G. Adelfio, "Efficient change point detection for genomic sequences of continuous measurements," *Bioinformatics*, vol. 27, no. 2, pp. 161–166, 2011.
- [4] L. Akoglu and C. Faloutsos, "Event detection in time series of mobile communication graphs," in *Proc. Army Sci. Conf.*, 2010, pp. 77–79.
- [5] K. Sequeira and M. Zaki, "ADMIT: Anomaly-based data mining for intrusions," in *Proc. ACM SIGKDD Int. Conf. on Knowledge Discovery and Data Mining*. ACM, 2002, pp. 386–395.
- [6] J. Mar, Y. C. Yeh, and I. F. Hsiao, "An ANFIS-IDS against deauthentication DoS attacks for a WLAN," in *Proc. IEEE Int. Symp. on Inform. Theory and Its Applications*, Oct 2010, pp. 548–553.
- [7] G. Androulidakis, V. Chatzigiannakis, S. Papavassiliou, M. Grammatikou, and V. Maglaris, "Understanding and evaluating the impact of sampling on anomaly detection techniques," in *Proc. IEEE Military Commun. Conf.*, Oct 2006, pp. 1–7.
- [8] G. Androulidakis and S. Papavassiliou, "Improving network anomaly detection via selective flow-based sampling," *IET Commun.*, vol. 2, no. 3, pp. 399–409, March 2008.
- [9] H. Wang, D. Zhang, and K. G. Shin, "Change-point monitoring for the detection of DoS attacks," *IEEE Trans. on Dependable and Secure Computing*, vol. 1, no. 4, pp. 193–208, Oct 2004.
- [10] R. J. Bolton and D. J. Hand, "Statistical fraud detection: A review," *Statistical Sci.*, pp. 235–249, 2002.
- [11] H. Sohn, J. A. Czarnecki, and C. R. Farrar, "Structural health monitoring using statistical process control," *J. Structural Eng.*, vol. 126, no. 11, pp. 1356–1363, 2000.
- [12] S. Xie, G. Wang, S. Lin, and P. S. Yu, "Review spam detection via temporal pattern discovery," in *Proc. ACM SIGKDD Int. Conf. on Knowledge Discovery and Data Mining*. ACM, 2012, pp. 823–831.
- [13] L. Lai, Y. Fan, and H. V. Poor, "Quickest detection in cognitive radio: A sequential change detection framework," in *Proc. IEEE Global Telecomm. Conf.*, Nov 2008, pp. 1–5.
- [14] N. Kundargi and A. Tewfik, "A nonparametric sequential Kolmogorov-Smirnov test for transmit opportunity detection at the MAC layer," in *Proc. IEEE Workshop on Signal Proc. Advances in Wireless Commun.*, June 2009, pp. 101–105.
- [15] F. Ibrahim and V. Sharma, "Novel distributed sequential nonparametric tests for spectrum sensing," in *Proc. IEEE Global Conf. Signal and Inform. Process.*, Dec 2014, pp. 1180–1184.
- [16] S. K. R. and V. Sharma, "Performance of spectrum sensing algorithms under fading, electromagnetic interference and outliers," in *Proc. Twenty First Nat. Conf. on Commun.*, Feb 2015, pp. 1–6.
- [17] Y. Zhang, W. P. Tay, K. H. Li, M. Esseghir, and D. Gaiti, "Learning temporal-spatial spectrum reuse," *IEEE Trans. Commun.*, vol. 64, no. 7, pp. 3092 – 3103, Jul. 2016.
- [18] E. Hanafi, P. A. Martin, P. J. Smith, and A. J. Coulson, "On the distribution of detection delay for quickest spectrum sensing," *IEEE Trans. on Commun.*, vol. 64, no. 2, pp. 502–510, Feb 2016.
- [19] M. X. Cheng, Y. Ling, and W. B. Wu, "In-band wormhole detection in wireless ad hoc networks using change point detection method," in *Proc. IEEE Int. Conf. on Commun.*, May 2016, pp. 1–6.
- [20] Z. Li, H. Lu, and Y.-P. Tan, "Video segmentation based on sequential change detection," in *Proc. IEEE Int. Conf. on Multimedia and Expo*, vol. 3, June 2004, pp. 1955–1958 Vol.3.
- [21] J. Tang, W. P. Tay, and Y. Wen, "Dynamic request redirection and elastic service scaling in cloud-centric media networks," *IEEE Trans. Multimedia*, vol. 16, no. 5, pp. 1434 – 1445, Aug. 2014.
- [22] X. Ren, K. H. Johansson, and L. Shi, "Quickest change detection with observation scheduling," *IEEE Trans. Autom. Control*, vol. 62, no. 6, pp. 2635–2647, June 2017.
- [23] M. D. Escobar and M. West, "Bayesian density estimation and inference using mixtures," *J. Amer. Stat. Assoc.*, vol. 90, no. 430, pp. 577–588, 1995.
- [24] J. Q. Li and A. R. Barron, "Mixture density estimation," in *Proc. Advances in Neural Inform. Process. Syst.*, 1999, pp. 279–285.
- [25] E. S. Page, "Continuous inspection schemes," *Biometrika*, vol. 41, no. 1/2, pp. 100–115, 1954.

- [26] G. Lorden, "Procedures for reacting to a change in distribution," *Ann. Math. Stat.*, pp. 1897–1908, 1971.
- [27] G. V. Moustakides, "Optimal stopping times for detecting changes in distributions," *Ann. Stat.*, pp. 1379–1387, 1986.
- [28] H. V. Poor and O. Hadjilias, *Quickest detection*. Cambridge University Press Cambridge, 2009, vol. 40.
- [29] A. Tartakovsky, I. Nikiforov, and M. Basseville, *Sequential analysis: Hypothesis testing and changepoint detection*. CRC Press, 2014.
- [30] V. V. Veeravalli and T. Banerjee, "Quickest change detection," *Academic press library in signal process.: Array and statistical signal process.*, vol. 3, pp. 209–256, 2013.
- [31] D. Siegmund and E. S. Venkatraman, "Using the generalized likelihood ratio statistic for sequential detection of a change-point," *Ann. Stat.*, pp. 255–271, 1995.
- [32] T. L. Lai, "Information bounds and quick detection of parameter changes in stochastic systems," *IEEE Trans. Inf. Theory*, vol. 44, no. 7, pp. 2917–2929, 1998.
- [33] T. Banerjee and V. V. Veeravalli, "Data-efficient minimax quickest change detection with composite post-change distribution," *IEEE Trans. Inf. Theory*, vol. 61, no. 9, pp. 5172–5184, 2015.
- [34] M. Pawlak and A. Steland, "Nonparametric sequential signal change detection under dependent noise," *IEEE Trans. Inf. Theory*, vol. 59, no. 6, pp. 3514–3531, June 2013.
- [35] E. Sakurai and K. Yamanishi, "Comparison of dynamic model selection with infinite HMM for statistical model change detection," in *Proc. IEEE Inform. Theory Workshop*, Sept 2012, pp. 302–306.
- [36] L. Gordon and M. Pollak, "An efficient sequential nonparametric scheme for detecting a change of distribution," *Ann. Stat.*, pp. 763–804, 1994.
- [37] D. M. Hawkins and Q. Deng, "A nonparametric change-point control chart," *J. Quality Technology*, vol. 42, no. 2, 2010.
- [38] E. Rafajlowicz, M. Pawlak, and A. Steland, "Nonparametric sequential change-point detection by a vertically trimmed box method," *IEEE Trans. Inf. Theory*, vol. 56, no. 7, pp. 3621–3634, July 2010.
- [39] —, "Nonparametric sequential change-point detection by a vertical regression method," in *Proc. IEEE Workshop on Statistical Signal Process.*, Aug 2009, pp. 614–617.
- [40] M. Lang and A. M. Zoubir, "A nonparametric cumulative sum scheme based on sequential ranks and adaptive control limits," in *Proc. European Signal Process. Conf.*, Aug 2015, pp. 1984–1988.
- [41] G. J. Ross, D. K. Tasoulis, and N. M. Adams, "Nonparametric monitoring of data streams for changes in location and scale," *Technometrics*, vol. 53, no. 4, pp. 379–389, 2011.
- [42] S. Li, Y. Xie, H. Dai, and L. Song, "M-statistic for kernel change-point detection," in *Proc. Advances Neural Inform. Process. Syst.*, 2015, pp. 3366–3374.
- [43] B. Darkhovskii and B. Brodskii, "A nonparametric method for fastest detection of a change in the mean of a random sequence," *Theory of Probability & Its Applicat.*, vol. 32, no. 4, pp. 640–648, 1988.
- [44] S. Nitinawarat and V. V. Veeravalli, "Universal quickest outlier detection and isolation," in *Proc. IEEE Int. Symp. Inform. Theory*. IEEE, 2015, pp. 770–774.
- [45] G. J. Ross and N. M. Adams, "Two nonparametric control charts for detecting arbitrary distribution changes," *J. Quality Technology*, vol. 44, no. 2, p. 102, 2012.
- [46] T. S. Lau, W. P. Tay, and V. V. Veeravalli, "Quickest change detection with unknown post-change distribution," in *Proc. IEEE Int. Conf. Acoust., Speech, and Signal Process.*, 2017, pp. 3924–3928.
- [47] G. Lorden and M. Pollak, "Sequential change-point detection procedures that are nearly optimal and computationally simple," *Sequential Analysis*, vol. 27, no. 4, pp. 476–512, 2008.
- [48] T. S. Lau, W. P. Tay, and V. V. Veeravalli, Supplementary materials for "A binning approach to quickest change detection with unknown post-change distribution". [Online]. Available: https://drive.google.com/open?id=1xQN4rGxRVZf4m-KBew_U1sKzfvHWJce3
- [49] J. R. Kwapisz, G. M. Weiss, and S. A. Moore, "Activity recognition using cell phone accelerometers," *ACM SigKDD Explorations Newsletter*, vol. 12, no. 2, pp. 74–82, 2011.
- [50] C. D. Fuh, "SPRT and CUSUM in hidden Markov models," *The Annals of Statistics*, vol. 31, no. 3, pp. 942–977, 2003.
- [51] —, "Asymptotic operating characteristics of an optimal change point detection in hidden Markov models," *The Annals of Statistics*, vol. 32, no. 5, pp. 2305–2339, 2004.
- [52] D. Siegmund, *Sequential Analysis: Tests and Confidence Intervals*. Springer Sci. & Bus. Media, 1985.
- [53] R. N. Bhattacharya and E. C. Waymire, *A basic course in probability theory*. Springer, 2007, vol. 69.
- [54] W. Hoeffding, "Probability inequalities for sums of bounded random variables," *J. of the Amer. statistical association*, vol. 58, no. 301, pp. 13–30, 1963.