

Sequential (Quickest) Change Detection: Classical Results and New Directions

Liyan Xie, *Student Member, IEEE*, Shaofeng Zou, *Member, IEEE*,
Yao Xie, *Member, IEEE*, and Venugopal V. Veeravalli, *Fellow, IEEE*

Abstract—Online detection of changes in stochastic systems, referred to as sequential change detection or quickest change detection, is an important research topic in statistics, signal processing, and information theory, and has a wide range of applications. This survey starts with the basics of sequential change detection, and then moves on to generalizations and extensions of sequential change detection theory and methods. We also discuss some new dimensions that emerge at the intersection of sequential change detection with other areas, along with a selection of modern applications and remarks on open questions.

I. INTRODUCTION

THE efficient detection of abrupt changes in the statistical behavior of streaming data is a classical and fundamental problem in signal processing and statistics. The abrupt change-point usually corresponds to a triggering event that could have catastrophic consequences if it is not detected in a timely manner. Therefore, the goal is to detect the change as quickly as possible, subject to false alarm constraints. Such problems have been studied under the theoretical framework of sequential (or quickest) change detection [160], [194], [215]. With an increasing availability of high-dimensional streaming data, sequential change detection has become a centerpiece for many real-world applications, ranging from monitoring power networks [37], internet traffic [100], cyber-physical systems [142], sensor networks [164], social networks [152], [165], epidemic detection [17], scientific imaging [162], genomic signal processing [179], seismology [7], video surveillance [109], and wireless communications [95].

In various applications, the streaming data is high-dimensional and collected over networks, such as social networks, sensor networks, and cyber-physical systems. For this

Manuscript received October 28, 2020; revised January 28, 2021; accepted April 5, 2021. V. V. Veeravalli was supported in part by the Army Research Laboratory under Cooperative Agreement W911NF-17-2-0196 (IoBT CRA), and in part by the National Science Foundation (NSF) under grant CCF-16-18658 and CIF-15-14245, through the University of Illinois at Urbana-Champaign. L. Xie and Y. Xie were partially supported by an NSF CAREER Award CCF-1650913, DMS-1938106, DMS-1830210, CCF-1442635, and CMMI-1917624. S. Zou was partially supported by an NSF Award CCF-1948165.

L. Xie and Y. Xie are with H. Milton Stewart School of Industrial and Systems Engineering, Georgia Institute of Technology, Atlanta, GA 30332 USA (email: lxie49@gatech.edu; yao.xie@isye.gatech.edu).

S. Zou is with the Department of Electrical Engineering, University at Buffalo, The State University of New York, Buffalo, NY 14260 USA (email: szou3@buffalo.edu).

V. V. Veeravalli is with the Department of Electrical and Computer Engineering, University of Illinois at Urbana-Champaign, Urbana, IL 61801 USA (email: vvv@illinois.edu).

reason, the modern sequential change detection problem's scope has been extended far beyond its traditional setting, often challenging the assumptions made by classical methods. These challenges include complex spatial and temporal dependence of the data streams, transient and dynamic changes, high-dimensionality, and structured changes, as explained below. These challenges have fostered new advances in sequential change detection theory and methods in recent years.

(1) *Complex data distributions.* In modern applications, sequential data could have a complex spatial and temporal dependency, for instance, induced by the network structure [16], [68], [167]. In social networks, dependencies are usually due to interaction and information diffusion [116]: users in the social network have behavior patterns influenced by their past, while at the same time, each user in the network will be influenced by friends and connections. In sensor networks for river contamination monitoring [34], sensor observations tend to be spatially and temporally correlated.

(2) *Data dynamics.* The statistical behavior of sequential data is often non-stationary, particularly in the post-change regime due to the dynamic behavior of the anomaly that causes the change. For example, after a linear outage in the power systems, the system's transient behavior is dominated by the generators' inertial response, and the post-change statistical behavior can be modeled using a sequence of temporally cascaded transient phases [171].

(3) *High-dimensionality.* Sequential data in modern applications is usually high-dimensional. For example, in sensor networks, the Long Beach 3D seismic array consists of approximately 5300 seismic sensors that record data continuously for seismic activity detection and analysis. Changes in high-dimensional time series usually exhibit low-dimensional structures in the form of sparsity, low-rankness, and subset structures, which can be exploited to enhance the capability to detect weak signals quickly.

In this tutorial, our aim is to introduce standard methods and fundamental results in sequential change detection, along with recent advances. We also present new dimensions at the intersection of sequential change detection with other areas, as well as a selection of modern applications. We should emphasize that our focus is on *sequential* change detection, where the goal is to detect the change from sequential data in *real-time* and *as soon as possible*. Another important line of related research is *offline* change detection (e.g., [59], [188]), where the goal is to identify and localize changes in data sequence in a *retrospective* manner, which is not our

focus here. Prior books and surveys on related topics include, for instance, change detection for dynamic systems [97], sequential analysis [98], [194], sequential change detection [19], [160], [215], Bayesian change detection [201], change detection assuming known pre- and post-change distributions [159] and using likelihood-based approaches [186], as well as time-series change detection [6].

The rest of the survey is organized as follows. In Section II, we present the basic problem setup and classical results. In Section III, we discuss several extensions and generalizations of the classical methods. In Section IV, we discuss new dimensions which intersect with sequential change detection, with some remarks on open questions. In Section V, we present some modern applications of sequential change detection. In Section VI, we make some concluding remarks.

II. CLASSICAL RESULTS

A. Problem Definition

In the sequential change detection problem, also known as the quickest change detection (QCD) problem [131], [160], [215], the aim is to detect a possible change in the data generating distribution of a sequence of observations $\{X_n, n = 1, 2, \dots\}$. The initial distribution of the observations is the one corresponding to normal system operation. At some unknown time γ (referred to as the *change-point*), due to some event, the distribution of the random observations changes. The goal is to detect the change as quickly as possible, subject to false-alarm constraints. We start by assuming that the observations are independent and identically distributed (i.i.d.) with probability density function (pdf) f_0 before and pdf f_1 after the change-point, respectively. We discuss generalizations to non-i.i.d. observations in Section III.

To motivate the design of algorithms for sequential change detection, we consider the example of detecting a change in the mean of the data generating distribution. In Fig. 1(a), we plot a sample path of observations that are distributed according to a normal distribution with zero mean and unit variance $\mathcal{N}(0, 1)$ before the change-point of 500, and $\mathcal{N}(0.1, 1)$ after the change-point. As can be seen in Fig. 1(a), such a small mean shift cannot be detected through manual inspection of the samples. In Fig. 1(b), we plot the evolution path of a sequential change detection procedure, the CUSUM algorithm (which is discussed in detail in Section II-C2), applied to the observations in Fig. 1(a). It can be seen that the test statistic stays close to zero before the change and has a positive drift after the change. Therefore, the change can be detected by comparing the test statistic to a positive threshold b (for instance, $b = 2$) and raising an alarm when the test statistic exceeds the threshold for the first time. For the sample path in Fig. 1(a), this approach incurs a detection delay of 60 samples (if we take samples daily, this means a detection delay of 2 months; if the sampling rate is 60 samples per second, this means a detection delay of one second). One natural question to ask is that: can we do better, at least on average? Clearly, if we set a lower threshold, for instance $b = 1$, we can detect the change much more quickly. However, this would result in a false alarm at $k = 112$. This example illustrates the

tradeoff between false-alarm and detection delay, which is a central problem when designing sequential change detection procedures. The goal in sequential change detection theory is to find detection procedures that have guaranteed optimality properties in terms of this tradeoff.

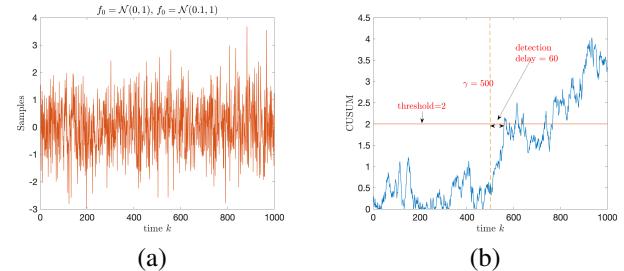


Fig. 1. To motivate the need for sequential change detection procedures, we plot a sample path with samples distributed according to $\mathcal{N}(0, 1)$ before the change and $\mathcal{N}(0.1, 1)$ after the change. We set the change-point $\gamma = 500$. As illustrated in (a), such a small mean shift cannot be detected through manual inspection of the samples. In (b), we plot the evolution of the CUSUM algorithm (detailed in Section II-C2) corresponding to the observations in (a), which can detect the change quickly.

B. Mathematical Preliminaries

Sequential change detection is closely related to the problem of statistical hypothesis testing, in which observations, whose distribution depends on the hypothesis, are used to decide which of the hypotheses is true. For the special case of binary hypothesis testing, we have two hypotheses, the *null hypothesis* and the *alternative hypothesis*. The classic Neyman-Pearson Lemma [136] establishes the form of the optimal test for this problem. In particular, consider the case of a single observation X , and suppose the pdf of X under the null and alternative hypotheses are f_0 and f_1 , respectively. Then, the test that minimizes the false negative error (Type-II error), under the constraint of the false positive error (Type-I error), is to compare the *likelihood ratio* $f_1(X)/f_0(X)$ to a threshold to decide which hypothesis is true. The likelihood ratio test is also optimal under other criteria such as Bayesian and minimax [131]. As we will see, the likelihood ratio also plays a key role in the development of sequential change detection algorithms.

The goal of sequential change detection is to design a *stopping time* on the observation sequence at which it is declared that a change has occurred. A stopping time is formally defined as follows:

Definition 1 (Stopping Time). *A stopping time with respect to a random sequence $\{X_n, n = 1, 2, \dots\}$ is a random variable τ such that for each n , the event $\{\tau = n\} \in \sigma(X_1, \dots, X_n)$, where $\sigma(X_1, \dots, X_n)$ denotes the sigma-algebra generated by (X_1, \dots, X_n) . Equivalently, the event $\{\tau = n\}$ is a function of only X_1, \dots, X_n .*

The main results on stopping times that are most useful for sequential change detection problems include Doob's Optional Stopping Theorem [43] and Wald's Identity [185].

A quantity that plays an important role in the performance of sequential change detection algorithms is the Kullback-Leibler (KL) divergence between two distributions.

Definition 2. (KL Divergence). The KL divergence between two pdfs f_1 and f_0 is defined as $D(f_1\|f_0) = \int f_1(x) \log(f_1(x)/f_0(x)) dx$.

Note that $D(f_1\|f_0) \geq 0$ with equality if and only if $f_1 = f_0$ almost surely. It is usually assumed that $0 < D(f_1\|f_0) < \infty$.

Define the log-likelihood ratio for an observation X :

$$\ell(X) := \log f_1(X)/f_0(X). \quad (1)$$

A fundamental property of the log-likelihood ratio, which is useful for constructing sequential change detection algorithms, is that before the change $n < \gamma$, the expected value of $\ell(X_n)$ is equal to $-D(f_0\|f_1) < 0$; and after the change, $n \geq \gamma$, the expected value of $\ell(X_n)$ is equal to $D(f_1\|f_0) > 0$. As will be seen later, the KL divergence between the pre- and post-change distributions is an important quantity that characterizes the tradeoff between the average detection delay and the false-alarm rate.

C. Common Sequential Change Detection Procedures

We now present several commonly used sequential change detection procedures, including the Shewhart chart, CUSUM, and Shiryaev-Roberts procedure, which enjoy certain optimality properties that we will make more precise later in Section II-D. These algorithms can be efficiently implemented in an online setting, which makes them useful in practice. We also briefly discuss some other sequential change detection procedures.

1) *Shewhart Chart*: One of the earliest sequential change detection procedures is the Shewhart chart [180], [181], which is widely used in industrial quality control [130]. The Shewhart chart was first introduced for the Gaussian model and based on comparing the instant observation to a threshold. We consider the log-likelihood-based modification and generalization of the standard Shewhart chart, where we compute the log-likelihood ratio based on the current observation (or the current batch of observations) and compare it with a threshold (called the control limit) to make a decision about the change. The property of the log-likelihood ratio discussed in Section II-B is utilized, which motivates the Shewhart chart:

$$\tau_{\text{sh}} = \inf \{n \geq 1 : \ell(X_n) > b\},$$

where b is a pre-specified threshold. The Shewhart chart is widely used in practice due to its simplicity. In [155], Pollak and Krieger showed that the Shewhart chart enjoys the optimality property that it maximizes the probability of detecting the change at the time it occurs, subject to false-alarm constraints. However, the Shewhart chart may suffer from “information loss” due to the fact that it ignores past observations in making a decision about the change, which leads to performance loss when we consider criteria, e.g., the tradeoff between detection delay and false-alarm rate (see Section II-D).

2) *Cumulative Sum (CUSUM) Procedure*: The CUSUM procedure, first introduced by Page [149], addresses the problem of “information loss” in the Shewhart chart. The CUSUM procedure uses past observations and thus can achieve a significant performance gain, especially when the change is small.

Although the CUSUM procedure was developed heuristically, it was later shown in [96], [122], [132], [170] that it has very strong optimality properties, which we will discuss further in Section II-D1.

The CUSUM procedure utilizes the properties of the cumulative log-likelihood ratio sequence:

$$S_n = \sum_{k=1}^n \ell(X_k).$$

Before the change occurs, the statistic has a negative drift because the expected value of $\ell(X_k)$ before the change is negative. After the change, it has a positive drift because the expected value of $\ell(X_k)$ after the change is positive. Thus, S_n roughly attains its minimum at the change-point γ . The CUSUM procedure is then constructed to detect this change in the drift of S_n . Specifically, the exceedance of S_n with respect to its past minimum is taken and compared with a threshold $b > 0$:

$$\tau_c = \inf \left\{ n \geq 1 : W_n = \left(S_n - \min_{0 \leq k \leq n} S_k \right) \geq b \right\}. \quad (2)$$

The CUSUM statistic can be rewritten as:

$$W_n = \max_{0 \leq k \leq n} \sum_{i=k+1}^n \ell(X_i) = \max_{1 \leq k \leq n+1} \sum_{i=k}^n \ell(X_i). \quad (3)$$

Note that the maximization over all possible $\gamma = k$ corresponds to plugging in a maximum likelihood estimate of the unknown change-point location in the log-likelihood ratio of the observations to form the CUSUM statistic. It can be shown that W_n can be computed recursively:

$$W_n = (W_{n-1} + \ell(X_n))^+, \quad W_0 = 0,$$

where $(x)^+ = \max\{x, 0\}$. This recursion enables the efficient online implementation of the CUSUM procedure in practice.

3) *Shiryaev-Roberts Procedure*: The maximum likelihood interpretation of the CUSUM procedure is closely related to another popular algorithm in the literature, called the Shiryaev-Roberts (SR) procedure. In the SR procedure, the maximum in (3) is replaced by a sum and the log-likelihood ratio is replaced by likelihood ratio. The detection statistic for the SR procedure is then defined as:

$$T_n := \sum_{1 \leq k \leq n} \prod_{i=k}^n e^{\ell(X_i)}, \quad (4)$$

and the corresponding stopping time is defined as

$$\tau_{\text{sr}} = \inf \{n \geq 1 : T_n \geq b\}.$$

The SR statistic can also be computed recursively:

$$T_n = (1 + T_{n-1}) e^{\ell(X_n)}, \quad T_0 = 0.$$

D. Optimality

We now briefly summarize optimality results in the existing literature for the above procedures. We begin by considering the non-Bayesian setting, where we do not assume a prior on the change-point γ , and then consider the Bayesian setting,

where the change-point is assumed to follow a certain distribution.

A fundamental problem in sequential change detection is to optimize the tradeoff between the false-alarm rate and the average detection delay, as illustrated in Section II-A using the example in Fig. 1. Controlling the false-alarm rate is commonly achieved by setting an appropriate threshold on a test statistic such as the one in (2). But the threshold also affects the average detection delay. A larger threshold incurs fewer false alarms but leads to a larger detection delay, and vice versa.

1) *Minimax Optimality*: In non-Bayesian settings, the change-point is assumed to be a *deterministic* and unknown variable. In this case, the average run length (ARL) to false alarm is generally used as a performance measure for false alarms:

$$\text{ARL}(\tau) = \mathbb{E}_\infty[\tau], \quad (5)$$

where \mathbb{P}_∞ is the probability measure on the sequence of observations when the change never occurs, and \mathbb{E}_∞ is the corresponding expectation. Its reciprocal, the false-alarm rate (FAR), is also commonly used:

$$\text{FAR}(\tau) = \frac{1}{\text{ARL}(\tau)} = \frac{1}{\mathbb{E}_\infty[\tau]}. \quad (6)$$

FAR can also be interpreted as the rate at which false alarms occur in the pre-change regime if we repeat the change detection procedure after each false alarm. Denote the set of stopping times that satisfy a constraint α on the FAR by:

$$\mathcal{D}_\alpha = \{\tau : \text{FAR}(\tau) \leq \alpha\}. \quad (7)$$

Finding a uniformly powerful test that minimizes the delay over all possible values of the change-point γ , subject to a FAR constraint, is generally intractable. Therefore, it is more tractable to pose the problem in the so-called minimax setting. There are two essential measures of the detection delay in the minimax setting, due to Lorden [122] and Pollak [154], respectively.

Lorden considers the supremum of the average detection delay conditioned on the worst possible realizations. In particular, Lorden defines¹:

$$\text{WADD}(\tau) = \sup_{n \geq 1} \text{ess sup } \mathbb{E}_n [(\tau - n)^+ | X_1, \dots, X_{n-1}], \quad (8)$$

where \mathbb{P}_n denotes the probability measure on the observations when the change occurs at time n , and \mathbb{E}_n denotes the corresponding expectation. We then have the following Lorden's formulation:

$$\text{minimize } \text{WADD}(\tau) \text{ subject to } \text{FAR}(\tau) \leq \alpha. \quad (9)$$

For the i.i.d. setting, Lorden showed that Page's CUSUM procedure given in (2) is asymptotically optimal as $\alpha \rightarrow 0$. It was later shown in [132] and [170] that a slight modification

¹Lorden defined WADD with $(\tau - n + 1)^+$ inside the expectation, i.e., he assumed a penalty of 1 if the algorithm stops at the change-point. We drop this additional penalty in our definition to be consistent with the other delay definitions in this paper.

of the CUSUM procedure, with $W_n = (W_{n-1})^+ + \ell(X_n)$, is exactly optimal for (9) for all $\alpha > 0$.

Although the CUSUM procedure is exactly optimal under Lorden's formulation, WADD(τ) is a pessimistic measure of detection delay since it considers the worst-case pre-change samples. An alternative measure of detection delay was suggested by Pollak [154]:

$$\text{CADD}(\tau) = \sup_{n \geq 1} \mathbb{E}_n[\tau - n | \tau \geq n], \quad (10)$$

for all stopping times τ for which the expectation is well-defined. It is easy to see that for any stopping time τ , $\text{WADD}(\tau) \geq \text{CADD}(\tau)$, and therefore, Pollak's formulation is less pessimistic.

In general, it may be challenging to exactly solve the problem in (9) and the corresponding problem defined using CADD in (10). For this reason, asymptotically optimal solutions for the above problems are often investigated in the literature. Specifically, a stopping time τ is said to be *first-order* asymptotically optimal if it satisfies:

$$\frac{\text{CADD}(\tau)}{\inf_{\tau \in \mathcal{D}_\alpha} \text{CADD}(\tau)} \rightarrow 1, \quad \text{as } \alpha \rightarrow 0;$$

it is *second-order* asymptotically optimal if $\text{CADD}(\tau)$ is within a constant of the best possible delay over the class \mathcal{D}_α :

$$\text{CADD}(\tau) - \inf_{\tau \in \mathcal{D}_\alpha} \text{CADD}(\tau) = O(1);$$

and it is *third-order* asymptotically optimal if such a constant goes to 0 as $\alpha \rightarrow 0$:

$$\text{CADD}(\tau) - \inf_{\tau \in \mathcal{D}_\alpha} \text{CADD}(\tau) = o(1).$$

These notions can also be defined similarly for the problem in (9) defined using WADD.

Pollak's formulation has been studied for the i.i.d. data in [154] and [197]. The first-order asymptotic optimality for Lorden's formulation can also be extended to Pollak's formulation. To show this, Lorden in [122] established a universal lower bound for WADD and Lai in [96] proved the lower bound to CADD:

Theorem 1 (Lower Bound for CADD [96]). As $\alpha \rightarrow 0$,

$$\inf_{\tau \in \mathcal{D}_\alpha} \text{CADD}(\tau) \geq \frac{|\log \alpha|}{D(f_1 || f_0)} (1 + o(1)).$$

It can be shown that the CUSUM procedure with a threshold $b = |\log \alpha|$ is first-order asymptotically optimum for both Lorden's and Pollak's formulations. In particular, as $\alpha \rightarrow 0$,

$$\text{CADD}(\tau_c) = \text{WADD}(\tau_c) \sim \frac{|\log \alpha|}{D(f_1 || f_0)},$$

where \sim means the ratio of the quantities on its two sides approaches 1 as $\alpha \rightarrow 0$.

The SR procedure is also asymptotically optimal and it was shown in [197] that by setting the threshold $b = 1/\alpha$,

$$\text{CADD}(\tau_{\text{SR}}) = \frac{|\log \alpha|}{D(f_1 || f_0)} + \xi + o(1),$$

where ξ is a constant that can be characterized using the nonlinear renewal theory [230] (details omitted here).

Finally, results in [133], [155], [196] show that the Shewhart chart is optimal for the criterion of maximizing the probability of detecting the change upon its occurrence subject to the FAR constraints. A more precise statement of this optimality property is as follows. Let the post-change density be denoted by $f_\theta(x)$, where $\theta \in \Theta$ is the post-change parameter. The Shewhart chart as defined earlier becomes the following stopping time:

$$\tau_{\text{Sh}} = \inf \left\{ n \geq 1 : \frac{f_\theta(X_n)}{f_0(X_n)} > b \right\},$$

where b is a pre-specified threshold. It is shown that when the threshold b is selected such that $\text{FAR}(\tau_{\text{Sh}}) = \alpha$, then τ_{Sh} is the optimal solution to the following optimization problem:

$$\text{maximize } \inf_{1 \leq n < \infty} \mathbb{P}_n^\theta(\tau = n | \tau \geq n) \text{ subject to } \text{FAR}(\tau) \leq \alpha, \quad (11)$$

where \mathbb{P}_n^θ denotes the probability when the change happens at n with θ being the post-change parameter. Moreover, it was shown in [196] that if the likelihood ratio $f_\theta(X)/f_0(X)$ is a monotone non-decreasing function of a statistic $S(X)$, then the Shewhart chart is equivalent to $\tau_{\text{Sh}} = \inf\{n \geq 1 : S(X_n) > b\}$ and when b is selected such that $\text{FAR}(\tau_{\text{Sh}}) = \alpha$, the Shewhart chart is uniform optimal in $\theta \in \Theta$ in the sense of solving (11) for all $\theta \in \Theta$.

In summary, both the CUSUM and SR procedures are asymptotically optimal with respect to Lorden's formulation and Pollak's formulation. The FAR decays to zero exponentially with exponent $D(f_1||f_0)$. We demonstrate the theory using an example in Fig. 2 by plotting the tradeoff curve between the CADD and $-\log(\text{FAR})$ for the CUSUM procedure. Note that the curve has a slope approximately of $1/D(f_1||f_0)$, which is consistent with the theory.

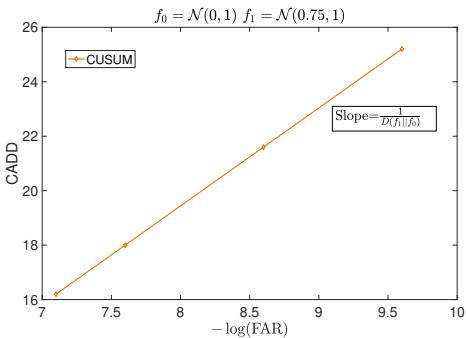


Fig. 2. Tradeoff curve between CADD and $-\log(\text{FAR})$ for the CUSUM algorithm. The pre-change distribution is $f_0 = \mathcal{N}(0, 1)$, and the post-change distribution is $f_1 = \mathcal{N}(0.75, 1)$. The slope of the curve is approximately $1/D(f_1||f_0)$.

Some more optimality results are summarized as follows. Under Pollak's criterion, it was shown in [197] that the SR algorithm is second-order asymptotically optimal, and that the SRP algorithm (Pollak's version of the SR algorithm that starts from a quasi-stationary distribution of the SR statistic) is third-order asymptotically optimal (as was also first established in [154]). More importantly, in [197], it was proved that the SR- r procedure that starts from a specially selected fixed point r is third-order optimal. In [158], it was shown that SR- r is strictly

optimal for CADD in some special cases. We also note that the (generalized) Shewhart chart is optimal for the criterion of maximizing the probability of detecting a change subject to false alarm constraints.

2) *Bayesian Optimality*: In the Bayesian setting, it is assumed that the change-point is a random variable Γ taking values on the non-negative integers, with probability mass function $\pi_n = \mathbb{P}\{\Gamma = n\}$. For a stopping time τ , define the average detection delay (ADD) and the probability of false alarm (PFA) as follows:

$$\text{ADD}(\tau) = \mathbb{E}[(\tau - \Gamma)^+] = \sum_{n=0}^{\infty} \pi_n \mathbb{E}_n[(\tau - \Gamma)^+] \quad (12)$$

$$\text{PFA}(\tau) = \mathbb{P}(\tau < \Gamma) = \sum_{n=0}^{\infty} \pi_n \mathbb{P}_n(\tau < \Gamma). \quad (13)$$

In Bayesian sequential change detection, the goal is to minimize ADD subject to a constraint on PFA. Shiryaev [183] formulated the Bayesian sequential change detection problem as follows:

$$\text{minimize } \text{ADD}(\tau) \text{ subject to } \text{PFA}(\tau) \leq \alpha. \quad (\text{Shiryaev}) \quad (14)$$

The prior on the change-point Γ is usually assumed to be a geometric distribution with parameter $0 < \rho < 1$,

$$\pi_n = \mathbb{P}\{\Gamma = n\} = \rho(1 - \rho)^{n-1} \mathbb{I}_{\{n \geq 1\}}, \quad \pi_0 = 0, \quad (15)$$

where \mathbb{I} is the indicator function. The justification for this assumption is that the geometric distribution is memoryless. Moreover, it leads to a tractable formulation and convenient optimal solutions to the Bayesian problem in (14) as we will discuss in the following.

The detection statistic of the *Shiryaev algorithm* is the posterior probability that the change has taken place given the observations so far. Denote by $X_1^n = (X_1, \dots, X_n)$ the observations up to time n , and by

$$p_n = \mathbb{P}(\Gamma \leq n | X_1^n) \quad (16)$$

the *a posteriori* probability at time n that the change has taken place given the observations up to time n . It then follows from the Bayes' rule that p_n can be updated recursively:

$$p_{n+1} = \frac{\tilde{p}_n e^{\ell(X_{n+1})}}{\tilde{p}_n e^{\ell(X_{n+1})} + (1 - \tilde{p}_n)}, \quad (17)$$

where $\tilde{p}_n = p_n + (1 - p_n)\rho$, and $p_0 = 0$. Then the *Shiryaev algorithm* is defined by comparing p_n with a given threshold b_α :

$$\tau_s = \inf \{n \geq 1 : p_n \geq b_\alpha\}, \quad (18)$$

where $b_\alpha \in (0, 1)$ is chosen such that the false alarm constraint, $\text{PFA}(\tau_s) \leq \alpha$, is satisfied.

Theorem 2 (Optimal Bayesian Procedure [183], [184]). *When the threshold b_α is selected such that $\text{PFA}(\tau_s) = \alpha$, the Shiryaev algorithm in (18) is Bayesian optimal for (14).*

An equivalent form of the Shiryaev statistic can be developed using the idea of the likelihood ratio test. This builds a connection to the earlier SR statistic defined in (4), and

it reveals useful insights about the nature of the procedure. Consider two hypotheses: “ $H_1 : \Gamma \leq n$ ” and “ $H_0 : \Gamma > n$ ”. Denote by $R_{n,\rho} = p_n/[\rho(1 - p_n)]$ the scaled likelihood ratio between the two hypotheses averaged over the change-point. It then follows that $R_{n,\rho}$ can be updated recursively as:

$$R_{n+1,\rho} = \frac{1 + R_{n,\rho} e^{\ell(X_{n+1})}}{1 - \rho}, \quad R_{0,\rho} = 0. \quad (19)$$

The Shiryaev stopping time τ_s in (18) can then be rewritten as a comparison of $R_{n,\rho}$ with a threshold. We remark here that if we set $\rho = 0$, then the Shiryaev statistic reduces to the SR statistic in (4).

A generalized Shewhart chart is also Bayesian optimal, as shown in [155], in the sense that it minimizes the expected loss where the loss function is $\mathbb{I}_{\{\tau \neq \Gamma\}}$, assuming that the change-point Γ follows a geometric prior and the parameter θ of the post-change distribution follows a known prior G . This result was generalized in [196, Theorem 5.1]. Moreover, both the CUSUM and SR procedures are first-order asymptotically optimal for the Bayesian setting when the prior has a heavy tail, or when the change-point is geometrically distributed with a small enough parameter.

3) *Evaluating the Performance Metrics*: In the definition of the WADD metric (8) and the CADD metric (10), it appears that we need to consider the supremum over all possible past observations and over all possible change-points. However, we can actually show that for the CUSUM and SR procedures, and some other algorithms, that the supremum over all possible change-points in WADD and CADD is achieved at time $n = 1$:

$$\begin{aligned} \text{CADD}(\tau_c) &= \text{WADD}(\tau_c) = \mathbb{E}_1 [\tau_c - 1], \\ \text{CADD}(\tau_{\text{SR}}) &= \text{WADD}(\tau_{\text{SR}}) = \mathbb{E}_1 [\tau_{\text{SR}} - 1]. \end{aligned}$$

Therefore, the CADD and the WADD can be conveniently evaluated by setting $\gamma = 1$, without “taking the supremum”.

E. Other Sequential Change Detection Procedures

1) *Mixture and Generalized Likelihood Ratio (GLR) Statistics*: The CUSUM and SR procedures require full knowledge of pre- and post-change distributions to obtain the log-likelihood ratio $\ell(X)$ used in computing the test statistics. In practice, the post-change distribution f_1 may be unknown. In the parametric setting, the post-change distribution can be parametrized using f_θ , where $\theta \in \Theta$ is the unknown parameter. Two commonly used methods for the situation here, which corresponds to the problem of composite hypothesis testing, are the generalized likelihood ratio (GLR) approach and the mixture approach. In the GLR approach, a supremum over $\theta \in \Theta$ is taken in constructing the test statistic. In particular, the test statistic for the GLR-CUSUM algorithm is given by:

$$W_n^G = \max_{1 \leq k \leq n+1} \sup_{\theta \in \Theta} \sum_{i=k}^n \ell_\theta(X_i), \quad (20)$$

where $\ell_\theta(X) = \log(f_\theta(X)/f_0(X))$. Performance analyses of the GLR-CUSUM algorithm for one-parameter exponential families can be found in [122], [123]. A major drawback of the GLR approach is that the corresponding GLR statistic (e.g., the one given in (20)) cannot be computed recursively in time,

except in some special cases (e.g., when the parameter set Θ has finite cardinality). To reduce the computational cost, a window-limited GLR approach was developed in [229] and generalized in [96], [99]. Window-limited versions of the GLR algorithm can be shown to be asymptotically optimal in certain cases if the window size is carefully chosen as a function of FAR.

The mixture method replaces the supremum over $\theta \in \Theta$ by a weighted average. For example, the mixture-CUSUM statistic is computed as:

$$W_n^m = \max_{1 \leq k \leq n+1} \log \int_{\Theta} \prod_{i=k}^n \frac{f_\theta(X_i)}{f_0(X_i)} \omega(\theta) d\theta, \quad (21)$$

where $\omega(\theta)$ is a weight function that integrates (sums) to 1 over Θ . Note that, like the GLR test statistic, the mixture test statistic cannot be computed recursively in general. It was shown in [196] that the mixture approach can result in first-order asymptotically optimal tests for practically any prior for both the i.i.d. and non-i.i.d. cases. In [189], the optimal prior was established such that the resultant mixture SR procedure is asymptotically optimal in a certain stronger sense.

2) *EWMA*: Note that the CUSUM and SR procedures can achieve a significant gain in performance when compared to the Shewhart chart by making use of past observations, i.e., CUSUM and SR have memory. The exponentially weighted moving average (EWMA) chart is another type of sequential change detection procedure that employs past observations. The EWMA detection statistic was originally defined as $Z_n = \lambda X_n + (1 - \lambda)Z_{n-1}$, where $\lambda \in (0, 1]$ is a pre-specified constant, with the aim to detect mean shift. The EWMA can be generalized to $Z_n = \lambda \ell(X_n) + (1 - \lambda)Z_{n-1}$ to detect shift in distribution, more generally. Thus, Z_n is a weighted moving average of all past information with weights decreasing exponentially in time. The EWMA chart is simple to implement and does not require any prior knowledge of the pre- and post-change distributions. A performance comparison of the EWMA chart and the CUSUM and SR procedures is given in [157].

III. GENERALIZATIONS AND EXTENSIONS

A. General Asymptotic Theory for Non-i.i.d. Data

There has been a considerable amount of effort to generalize the optimality results for sequential change detection to the non-i.i.d. setting. Lai [96] initiated the development of a general minimax asymptotic theory for both Lorden's and Pollak's formulations, while Tartakovsky and Veeravalli [206] initiated the development of a general Bayesian asymptotic theory.

1) *General Minimax Asymptotic Theory*: Under the minimax setting, Lai in [96] obtained a general lower bound for non-i.i.d. data on the CADD (and hence on the WADD) for any stopping time that satisfies the constraint that FAR is no larger than α . It was then shown that an extension of the CUSUM procedure (2) to the non-i.i.d. setting achieves this lower bound asymptotically as $\alpha \rightarrow 0$. There are also works investigating non-i.i.d. data under some specific settings, e.g., multi-sensor slope change detection [28], linear regression models [63],

[221], generalized autoregressive conditional heteroskedasticity (GARCH) models [22], non-stationary time series [42], general stochastic models [195], [200], and hidden Markov models [62]. We refer to [196] for more recent developments on this topic.

We now present a generalized CUSUM procedure for non-i.i.d. data. In this setting, conditional distributions are used to compute the likelihood ratios. In the pre- and post-change regimes, the conditional distribution of X_n given X_1^{n-1} is given by $f_{0,n}(X_n|X_1^{n-1})$ and $f_{1,n}(X_n|X_1^{n-1})$, respectively. Define the conditional log-likelihood ratio and the CUSUM statistic, respectively, as:

$$Y_i = \log \frac{f_{1,i}(X_i|X_1^{i-1})}{f_{0,i}(X_i|X_1^{i-1})}, \quad \text{and } C_n = \max_{1 \leq k \leq n+1} \sum_{i=k}^n Y_i.$$

Then the stopping time for the generalized CUSUM is defined as:

$$\tau_G = \inf \{n \geq 1 : C_n \geq b\}. \quad (22)$$

Note the generalized CUSUM (for non-i.i.d. data) takes a similar form as the original CUSUM (for i.i.d. data) except that we replace the log-likelihood ratio with the conditional log-likelihood ratio.

The minimax optimality of the generalized CUSUM for the non-i.i.d. data was established in [96]. Under some regularity conditions, by setting the threshold $b = |\log \alpha|$, we have $\tau_G \in \mathcal{D}_\alpha$. If there exists I such that

$$\max_{m \leq t} \frac{1}{t} \sum_{i=n}^{n+m} Y_i \rightarrow I \quad \text{a.s. } \mathbb{P}_n, \quad \text{as } t \rightarrow \infty \quad \forall n, \quad (23)$$

and the convergence is complete in the sense that $\sum_{n=1}^{\infty} \mathbb{P}_1(|(1/n) \sum_{i=1}^n Y_i - I| \geq \epsilon) < \infty$ for all $\epsilon > 0$, then as $\alpha \rightarrow 0$:

$$\begin{aligned} \text{CADD}(\tau_G) &\sim \text{WADD}(\tau_G) \sim \inf_{\tau \in \mathcal{D}_\alpha} \text{WADD}(\tau) \\ &\sim \inf_{\tau \in \mathcal{D}_\alpha} \text{CADD}(\tau) \sim \frac{|\log \alpha|}{I}, \end{aligned} \quad (24)$$

where the positive constant $I > 0$ plays a similar role as the KL divergence in the i.i.d. setting.

2) *General Bayesian Asymptotic Theory*: Under the Bayesian setting, when the samples conditioned on the change-point are non-i.i.d., it is generally difficult to find an exact solution to the Shiryaev problem in (14). Tartakovsky and Veeravalli [206] showed that the Shiryaev algorithm is asymptotically optimal as $\alpha \rightarrow 0$, under some regularity conditions on the pre- and post-change distributions.

Similar to the i.i.d. case, we can define the posterior probability p_n of change having occurred before time n given all previous samples, in the same expression as (16). The Shiryaev algorithm for the non-i.i.d. setting is then defined in the same way as in (18). Note that the recursion in (17) may not hold for a general distribution for Γ . However, if the change-point Γ is geometrically distributed, a recursive expression for p_n can still be derived. Define

$$d = - \lim_{n \rightarrow \infty} \frac{\log \mathbb{P}(\Gamma > n)}{n},$$

which captures the decay rate of the tail probability of change-point Γ 's prior distribution as the sample size n increases. When Γ is “heavy-tailed”, $d = 0$, and when Γ has an “exponential tail”, $d > 0$. For example, when the prior distribution is geometric with parameter ρ as defined in (15), $d = |\log(1 - \rho)|$. If there exists I such that

$$\frac{1}{t} \sum_{i=n}^{n+t} Y_i \rightarrow I \quad \text{a.s. } \mathbb{P}_n \quad \text{as } t \rightarrow \infty, \quad \forall n, \quad (25)$$

and some additional conditions on the rates of convergence are satisfied (see [206] for the details), then the Shiryaev algorithm in (18) with a threshold $b_\alpha = 1 - \alpha$ is asymptotically optimal for Bayesian optimization problem in (14) as $\alpha \rightarrow 0$ [206]:

$$\text{ADD}(\tau_s) \sim \inf_{\tau: \text{PFA}(\tau) \leq \alpha} \text{ADD}(\tau) \sim \frac{|\log \alpha|}{I + d}. \quad (26)$$

Note that in [206], a general result for the m -th moment of the delay was developed. Here, for simplicity, we only presented the result for $m = 1$.

B. Change-of-measure to Obtain Accurate ARL Approximations

For CUSUM and SR procedures with i.i.d. samples, it may be relatively easy to evaluate their performance (such as the ARL) both theoretically and numerically, as discussed in Section II-D3. However, in many settings such as those involving non-i.i.d. observations, GLR statistics [188], and non-parametric statistics [115], it may be challenging to develop exact analytical expressions for the ARL (or its inverse the FAR). In these situations, one has to use onerous numerical simulation to obtain a threshold for a target ARL. To tackle this problem, techniques based on extremes in random fields have been developed [238], from which one can obtain accurate approximations to the ARL for many problems.

1) *Using Change-of-measure to Analyze the ARL*: The main idea here is to relate finding ARL to finding the tail probability of the maximum of a random field. To obtain a more accurate approximation of the ARL, an alternative probability measure is considered, under which false alarms are more likely to occur. This is analogous to “importance sampling”, but it is more involved since the alternative probability measure is usually a mixture of distributions.

The analysis usually involves two steps. First, we aim to find the probability $\mathbb{P}_\infty\{\tau \leq m\}$, for a large constant $m > 0$ and stopping time τ (the first time the detection statistic exceeds the threshold b). Finding this probability is challenging because $\{\tau \leq m\}$ is a rare event under the pre-change regime, especially when the threshold b is large (this is the asymptotic scenario that we are interested in). Therefore, the change-of-measure technique plays an important role by considering an alternative measure under which $\{\tau \leq m\}$ happens with a much higher probability. More specifically, we choose the alternative measure such that the expectation of the detection statistic equals the threshold b . Then, using the local central limit theorem and the local behavior of the correlated random field, we can obtain an analytical expression for the probability of $\{\tau \leq m\}$ under the alternative measure. The

probability under the alternative measure is then converted back to the probability under the original measure through Mill's ratio. The rigorous mathematical derivations can be found in [238].

Second, we will relate the above probability to the ARL, leveraging the fact that the stopping time τ as threshold $b \rightarrow \infty$ is asymptotically exponentially distributed [4], [187]. Although this fact only holds strictly for stopping times for algorithms such as the CUSUM and SR when observations are i.i.d. [156], this method has been widely used and is verified to be highly accurate in practice (see examples in [28], [115], [188], [236]). Thus, for a large m , $\mathbb{P}_\infty\{\tau \leq m\} \sim 1 - e^{-\lambda_b m}$, where λ_b is the parameter of the exponential distribution. By definition, the mean of the exponential distribution is $1/\lambda_b$, which corresponds to the ARL.

2) *Example: Analyzing MMD-based Sequential Change Detection Procedure:* Below, we illustrate the change-of-measure technique by analyzing the non-parametric kernel-based maximum mean discrepancy (MMD) statistics (details can be found in [115]). The kernel MMD divergence, which measures the distance between two arbitrary distributions, is widely adopted in signal processing and machine learning. Given two sets of samples $X := \{x_1, \dots, x_n\}$ generated i.i.d. from a distribution f_0 and $Y := \{y_1, \dots, y_n\}$ generated i.i.d. from a distribution f_1 , an unbiased estimator of the MMD between f_0 and f_1 is the following:

$$\text{MMD}(X, Y) = \frac{1}{n(n-1)} \sum_{i \neq j} \{k(x_i, x_j) + k(y_i, y_j) - k(x_i, y_j) - k(x_j, y_i)\},$$

where $k(\cdot, \cdot)$ is a kernel in the reproducing kernel Hilbert space (RKHS), e.g., Gaussian kernel. Intuitively, the MMD statistic is small when f_0 is similar to f_1 , and is large otherwise.

The sequential change detection procedure based on the MMD statistic is then defined as follows [115]. At each time t , we treat the most recent B samples, denoted by $X_{t-B+1}^t := \{X_{t-B+1}, \dots, X_t\}$, as the test block ($B > 0$ is a pre-specified parameter). Then we sample N blocks of size B from the ‘‘reference’’ data generated from the pre-change distribution, denoted by $\{\tilde{X}_1, \dots, \tilde{X}_N\}$. We compute an average MMD statistic of all the reference blocks with respect to the test block:

$$U_t = \frac{1}{N} \sum_{i=1}^N \text{MMD}(\tilde{X}_i, X_{t-B+1}^t).$$

Define $Z'_t := U_t / \sqrt{\text{var}[U_t]}$ as the standardized detection statistic, where the variance $\text{var}[U_t]$ can be found in closed-form and can be estimated conveniently from data [115]. The MMD-based procedure stops when the standardized MMD statistic exceeds a threshold b :

$$\tau_M = \inf\{t : Z'_t > b\}.$$

This corresponds to a generalized type of Shewhart chart.

Theorem 3 (ARL of MMD-based Procedure [115]). *Let $B > 0$. When $b \rightarrow \infty$, the ARL of the stopping time τ_M , $\mathbb{E}_\infty[\tau_M]$, is given by:*

$$\frac{e^{b^2/2}}{b} \left\{ \frac{2B-1}{\sqrt{2\pi B(B-1)}} \nu \left(b \sqrt{\frac{2(2B-1)}{B(B-1)}} \right) \right\}^{-1} (1 + o(1)),$$

where $\nu(\cdot)$ is a special function whose definition can be found in [185].

We present the main step of the proof to Theorem 3 to illustrate the change-of-measure technique. First, note that the event $\{\tau \leq m\}$ is the same as the maximum of the detection statistic has exceed the threshold b at some point before m , i.e., $\{\sup_{2 \leq t \leq m} Z'_t \geq b\}$, and

$$\begin{aligned} & \mathbb{P}_\infty \left\{ \sup_{2 \leq t \leq m} Z'_t \geq b \right\} \\ &= \mathbb{E}_\infty \left\{ \frac{\sum_{t=2}^m e^{\xi_t}}{\sum_{s=2}^m e^{\xi_s}} \mathbb{I}_{\{\sup_{2 \leq t \leq m} Z'_t \geq b\}} \right\} \\ &= e^{-b^2/2} \sum_{t=2}^m \mathbb{E}_t \{ R_t e^{-[\xi_t - b^2/2 + \log M_t]} \mathbb{I}_{\{\xi_t - b^2/2 + \log M_t \geq 0\}} \}, \end{aligned}$$

where $\xi_t = bZ'_t - b^2/2$ is the log-likelihood ratio between the changed measure $\mathbb{E}_t[X] = \mathbb{E}_\infty[X e^{\xi_t}]$ and the original measure \mathbb{E}_∞ , $M_t = \max_s e^{\xi_s - \xi_t}$, $S_t = \sum_s e^{\xi_s - \xi_t}$, and the so-called Mill's ratio $R_t = M_t / S_t$. The result in Theorem 3 is established by establishing properties of the local field $\{\xi_s - \xi_t\}$ and the global term $\xi_t - b^2/2$ (details omitted here).

The numerical example in Fig. 3 demonstrates that the threshold b (to achieve a target ARL) obtained using the theoretical approximation in Theorem 3 is consistent with that obtained from simulations, especially after a skewness correction. This example demonstrates that the theoretical approximation of the ARL obtained using the change-of-measure technique is of high accuracy, and thus can help avoid computationally expensive simulations to calibrate the procedure.

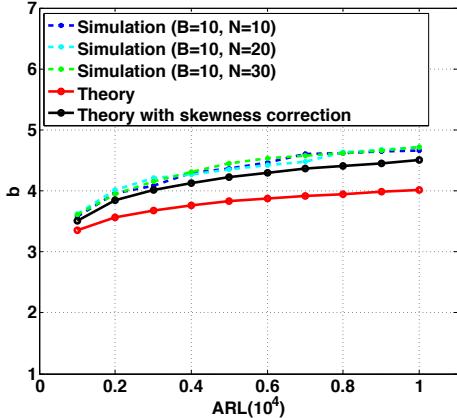


Fig. 3. Accuracy of ARL approximations, obtained by ‘‘change-of-measure’’, for the sequential MMD-based procedure: comparison of the thresholds obtained by simulation and from Theorem 3.

C. Non-stationary and Multiple Changes

In various modern applications, for instance, line outage detection in power systems [171] and stochastic power supply control in data centers [173], the change is not stationary. There can be a sequence of multiple changes: one followed by another. Below, we review some recent advances in sequential detection of dynamic changes.

1) Sequential Change Detection Under Transient Dynamics:

In classical sequential change detection formulations [19], [160], [194], [215], the statistical behavior of the observations is characterized by one pre-change distribution and *one* post-change distribution (known or unknown). In other words, the statistical behavior after the change is stationary. This assumption may be too restrictive for many practical applications with more involved statistical behavior after the change-point.

An example of the problem where the observations are non-stationary after the change, is sequential change detection under transient dynamics, which was studied in [171], [173], [174], [251]. Specifically, the pre-change distribution does not change to a persistent post-change distribution instantaneously, but after several transient phases, each phase is associated with a distinct data generating distribution. The goal is to detect the change as quickly as possible, either during the transient phases or during the persistent phase. This problem is fundamentally different from detecting a transient change (see, e.g., [51], [52], [64]), where the system goes back to the pre-change mode after a single transient phase, and the goal is to detect the change within the transient phase. The problem is also related to sequential change detection in the presence of a nuisance change, where the presence of the nuisance change can be modeled as a transient phase. However, an alarm should be raised only if the critical change occurs [103].

Two algorithms were proposed and investigated in [171], [251] for the minimax setting, the dynamic-CUSUM (D-CUSUM), and the weighted dynamic-CUSUM (WD-CUSUM), where the change-point and the transient durations are assumed to be unknown and deterministic. The basic idea is to construct a generalized likelihood based algorithm taking the supremum over the unknown change-point and the durations of transient phases. It was shown in [171], [251] that the D-CUSUM and WD-CUSUM test statistics can be updated recursively, and thus are computationally efficient. In [251], it was demonstrated that both algorithms are adaptive to the unknown transient dynamics, although durations of transient phases were unknown and were not employed in algorithm implementation. Moreover, both the D-CUSUM (under certain conditions) and the WD-CUSUM algorithms were shown to be first-order asymptotically optimal in [251]. The Bayesian setting was investigated in [174], where the change-point and the *durations* of transient phases are assumed to be geometrically distributed. The optimal test was constructed, and a computationally efficient alternative test based on thresholding the posterior probability that the change has occurred was also proposed.

2) *Sequential Detection of Moving Anomaly*: Existing studies on sequential change detection in networks usually assume that the change is persistent once it affects a node. However, there are scenarios where the change may not necessarily

be persistent at a particular node; instead, it is persistent across the network as a whole, e.g., a moving anomaly in a sensor network. In this case, existing approaches using CUSUM statistics from each node, e.g., [55], [66], [126], [255], cannot be applied. Recently, the problem of sequential moving anomaly detection in networks was studied in [175], [176]. Specifically, after an anomaly emerges in the network, one node is affected by the anomaly at each time instant and receives data from a post-change distribution. The anomaly dynamically moves across the network with an unknown trajectory, and the node that it affects changes with time. Two approaches have been proposed to model the trajectory of the anomaly: the hidden Markov model [176], and the worst-case approach [175], which we discuss in the following.

The first approach (hidden Markov model) [176] models the anomaly's trajectory as a Markov chain, and thus the samples are generated according to a hidden Markov model. The advantage of this model is that it takes into consideration the network's topology, i.e., that the anomaly only moves from a node to one of its neighbors. In [176], a windowed GLR based algorithm was constructed and was shown to be first-order asymptotically optimal. Alternative algorithms were also designed with performance guarantees, including the dynamic SR procedure, recursive change-point estimation, and a mixture CUSUM algorithm.

The second approach (worst-case approach) [175] assumes that the anomaly's trajectory is unknown but deterministic and considers the worst-case performance over all possible trajectories. A CUSUM-type procedure was constructed. The main idea is to use the mixture likelihood to construct a test statistic, which is further used to build a procedure of the CUSUM-type. This procedure was shown to be exactly optimal in [175] when the sensors are homogeneous. This idea has been further generalized to solve the sequential moving anomaly detection problem with heterogeneous sensors and has been shown to be first-order asymptotically optimal [172].

3) *Multiple Change Detection*: A related line of research is multiple change detection in the offline setting, which aims to estimate multiple change-points from observations in a retrospective study. Various methods were proposed to estimate the number and locations of change-points, including hierarchical clustering based method [125], binary segmentation type methods [9], [40], [41], [60], [61], [220], (penalized) least-squared methods [23], [106]–[108], [240], Schwarz criterion [239], kernel-based algorithms [8], [69], and so on. Another line of work aims to reduce the computational complexity of the multiple change detection methods, such as [71], [89], [169]. We refer to [210] for a recent review on multiple change detection. Some offline multiple change detection algorithms can motivate the development of their online versions.

4) *Decentralized and Asynchronous Change Detection in Networks*: When the information for detection is distributed across a network of sensors, detection problems fall under the umbrella of distributed (or decentralized) detection [31], [212], [214], [216]. In the decentralized setting, each sensor sends messages to the fusion center based on the observations it has received so far. The fusion center may provide feedback to sensors and make the final decision. The problem of

decentralized sequential change detection in distributed sensor systems was introduced in [217], considering the observation model where all sensors are affected by the change at the same time. There have been a number of papers on the topic since then, see e.g., [127], [205], [207]. A more recent (and practical) perspective is that the change may affect sensors with delay, i.e., different sensors may observe the change at different times, which we will present in the following.

In the case of multiple data streams, the change may happen asynchronously for different sensors. When we desire to detect the first onset of change, it is proposed in [66] to monitor each data stream by *local* CUSUM procedures and raise the alarm when any sensor raises an alarm. The sum of local CUSUM statistics has been considered in [126] and was shown to be asymptotically optimal. The problem where the change propagates from one sensor to the next with known Markov dynamics after the change was studied in [164], and an asymptotically optimal test was developed. A recent procedure proposed in [231] finds an optimal combination of local data streams accounting for their delays in being affected by the change, which can boost the signal-to-noise ratio and reduce the detection delay especially when the signal is weak.

In [255], the problem of sequentially detecting a significant change (i.e., when at least η number of sensors are affected by the change) was investigated. The event is dynamic, i.e., different nodes are affected at different times. Instead of using a scan statistic, which is computationally costly, a spartan-CUSUM (S-CUSUM) algorithm was constructed, which compares the sum of the smallest $N - \eta + 1$ local CUSUM statistics to a threshold, where N is the total number of nodes. For the case where the change propagates along network edges, a network-CUSUM (N-CUSUM) algorithm was further constructed based on the idea that the affected nodes shall induce a connected subgraph. The N-CUSUM algorithm was also shown to be first-order asymptotically optimal, and performs much better than the S-CUSUM numerically. The decentralized setting where there is no fusion center and nodes can only communicate with their neighbors was studied in [111], [253], and the approach is based on a novel combination of the alternating direction method of multipliers (ADMM) and average consensus approaches. In [94], a Bayesian approach is used to model the dynamic change with an unknown propagation pattern, where the goal is to detect the change when it firstly emerges in the network; an optimal solution structure is derived using a dynamic programming framework.

D. Robust Sequential Change Detection

Many classical procedures (for instance, CUSUM and SR) require exact knowledge of the pre- and post-change distributions. However, in real-world scenarios, the actual data distributions may be complex and different from what we have assumed. There can be adversarial attacks that significantly perturb the data distributions. This can lead to performance degradation of the optimal procedures. How to make the procedures more robust in the presence of model mismatch is the topic of robust sequential change detection.

1) Robustness to Model Uncertainties: There have been many efforts to make the detection procedure more robust to model uncertainties. One approach is to treat the pre- and post-change distributions to belong to some parametric family with unknown parameters in *uncertainty sets* and then form the GLR based test as we discussed earlier in Section II-E. Another approach to developing good tests in the presence of model uncertainties is through the use of minimax robustness as the criterion as is done in the seminal work of Huber on robust hypothesis testing [82], [83]. The solution to the robust hypothesis testing problem usually relies on finding the least favorable distributions (LFDs) within the uncertainty classes, with likelihood ratio of these distributions used in constructing the robust tests. It can be shown that LFDs exist for uncertainty classes satisfying a certain joint stochastic boundedness (JSB) condition [218]. The problem of minimax robust sequential change detection was explored in [213], in which an exactly optimal solution was obtained for uncertainty classes satisfying the JSB condition under a generalized Lorden criterion. An extension of this result to asymptotic minimax robust sequential change detection is studied in [129], where a weaker notion of stochastic boundedness is introduced.

A robust CUSUM algorithm is developed in [27] by making a connection to convex optimization, which is particularly useful for the high-dimensional setting and leads to a tractable formulation. For instance, assuming the covariance matrix lies in an *uncertainty set* centered around a nominal value, the problem of finding LFDs can be cast as solving a semidefinite program and can be solved efficiently.

2) Robustness to Adversarial Attacks: The problem of sequential change detection in sensor networks in the presence of adversarial attacks [102] was investigated in [20], [56]. In the presence of Byzantine attacks, an adversary may modify observations arbitrarily to defer the detection of a change and increase the false alarm rate. In [20], it is assumed that the change affects all but one compromised sensor, and the detection strategy is to raise a global alarm until two local CUSUMs exceed the threshold. In [56], a more general setting was investigated, where an unknown subset of sensors can be compromised. Sequential detection strategies were designed by waiting until L local CUSUM statistics exceed the threshold (simultaneously or not) or by comparing the sum of the L smallest CUSUM statistics to a threshold. With a proper choice of L , the above approaches are robust to Byzantine attacks.

E. Data-efficient Sequential Change Detection

There is usually a cost associated with making observations in practical engineering applications, e.g., the power consumption in sensor networks. An extension of Shiryaev's formulation (Section II-D2) was investigated in [11] by including an additional constraint on the average number of observations taken before the change. The cost of observations after the change is included in the detection delay. Specifically, whether to take an observation at time t is controlled by an on-off binary control variable S_t , and S_t is a function of all the information available up to time $t - 1$. A data-efficient

Shiryayev (DE-Shiryayev) algorithm was constructed in [11], and was shown to be asymptotically optimal as PFA goes to zero. The DE-Shiryayev algorithm is also shown to have good observation cost-delay tradeoff curves: for moderate values of PFA, for Gaussian observations, the delay of the algorithm is within 10% of the Shiryaev delay even when the observation cost is reduced by more than 50%. Furthermore, the DE-Shiryayev algorithm is substantially better than the standard approach of *fractional sampling* scheme, where the Shiryaev algorithm is used and where the observations to be skipped are determined *a priori* in order to meet the observation constraint. A minimax formulation was further proposed in [13] to address the scenario when a prior on the change-point is not available. The DE-CUSUM algorithm developed in [13] is shown to be asymptotically optimal as FAR goes to zero, and significantly outperforms fractional sampling in simulations. Extensions to composite post-change distributions were studied in [14], and generalizations to distributed sensor networks were explored in [15].

F. High-dimensional Streaming Data

High-dimensional data usually have low-dimensional structures, such as sparsity and low-rankness, which can be leveraged to achieve improved detection performance and computational efficiency. Meanwhile, missing data is very common for high-dimensional streaming data. In this section, we review recent advances in these directions.

1) *Sparse Change in Multiple Data Streams*: For multiple independent streams of data, a mixture procedure was developed in [236] to monitor parallel streams for a change-point that affects only a subset of them (usually sparse). Both the subset being affected and the post-change distribution are unknown. The mixture model hypothesizes that each sensor is affected with a small probability $\varrho \in (0, 1)$ by the change, where ϱ is pre-specified. The mixture detection statistic at time t is defined as

$$\sum_{n=1}^N \log \left[1 - \varrho + \varrho f_1(X_t^{(n)}) / f_0(X_t^{(n)}) \right],$$

where $X_t^{(n)}$ denotes the observation at the n -th sensor and at time t , and N is the number of sensors. Another efficient global monitoring scheme was proposed in [227] by combining hard thresholding with linear shrinkage estimator for the post-change parameters. In recent works [121], [247], a similar problem was tackled by running local detection procedures and using the sum of the shrinkage transformation of local detection statistics as a global detection statistic. This sum-shrinkage framework was further extended in [248] to be more robust to outliers using the Box-Cox transformation. Recent work [53] studied change detection in regimes where the dimension tends to infinity and the length of the sequence grows with the dimension.

2) *Subspace Change Detection*: In many applications, the change in high-dimensional data covariance structure can be represented as a low-rank change. For instance, in seismic signal detection [232], a similar waveform is observed at a subset of sensors after the change. Such a change can be

modeled as the covariance matrix shifts from an identity matrix to a “spiked” covariance model [88]. The subspace-CUSUM procedure was developed in [232], in which the *unknown* subspace in the post-change spiked model is estimated sequentially and further used to obtain the log-likelihood ratio statistic. A CUSUM procedure for detecting switching subspace (from a known subspace to another target subspace) was studied in [86].

3) *Missing Data*: In high-dimensional time series, it is common that we cannot observe all the entries at each time. The missing components in the observed data handicap conventional approaches. In [234], a mixture type of approach was proposed by combining subspace tracking with missing data to model the underlying dynamic of data geometry (submanifold). Specifically, streaming data is used to track a submanifold approximation, to measure deviations from this approximation, and to calculate a series of statistics of the deviations for detecting when the underlying manifold has changed.

4) *Sketching to Conquer High-dimensionality*: To detect changes quickly over high-dimensional data, we may need to conquer the challenges presented by the data’s high dimensionality. *Sketching* is a commonly used strategy to reduce data dimensionality, which performs linear projections of high-dimensional data into a small number of sketches. A GLR procedure based on data sketches was studied in [237], with the precise characterization of performance metrics and the minimum number of sketches needed to achieve good performance. Multiple types of sketching matrices can be used, such as Gaussian random matrices, expander graphs, and network topology constrained matrices. The sketching procedure is relevant to large power networks where we cannot place a sensor on each node or edge. Instead, each sensor will measure *aggregates* of the network states at a few edges or nodes. In [237], the mean-shift detection problem in power networks is studied, where each measurement corresponds to a linear combination of the state at an edge, e.g., real power flow. This leads to a sketching matrix determined by the network topology.

G. Joint Detection and Estimation

It is common that the distribution after the change is unknown. For instance, before the change in industrial process monitoring applications, the production line is in-control and well-calibrated (thus the distribution before the change is known). However, after the change, an anomaly causes a shift to the operation into an unknown status. Therefore, it is interesting to incorporate estimates of the possible post-change status into the detection statistic when performing detection; this problem is related to robust sequential change detection, as discussed in Section III-D1. In other situations, we need to estimate the post-change distribution in retrospect for identifying the change. There has been much work establishing the theoretical foundation for joint detection and estimation. For instance, [135] combines the Bayesian formulation of the estimation and detection and develops an optimal procedure to achieve a tradeoff between detection power and estimation

quality. In another context, it is also referred to as sequential change diagnosis [46]. Quickest *searching* of the change-point (e.g., quickest search for rare events) has been developed in [78], [79], [193].

H. Spatio-temporal Change Detection

When modeling discrete event data, the point process model [45] is frequently used due to its capability of modeling the time intervals between events directly. Point processes assume that time intervals between events are exponentially distributed. For example, in Poisson processes the intervals are independent, and in Hawkes processes the intervals are dependent, and the intensity depends on the events that occurred in the past [54]. The “autoregressive” nature of Hawkes processes makes them attractive in modeling temporal dependence and causal relationships, including market models [209], earthquake event prediction [144], inferring leadership in e-mail networks [57], and topic models [75]. The multi-dimensional Hawkes process model over networks can model highly correlated discrete event data [168] and capture dependence over networks and propagation of the signal in such settings.

Detection of changes for point processes has attracted much attention for both single event stream and multiple streams over networks (or over multiple locations). For example, there are works focusing on Poisson processes [76], [179], [246], and some recent work on one-dimensional [124], [153] and multi-dimensional (network) point processes [116], [226]. In particular, [116] studied the change detection for networked streaming event data and constructed GLR type procedures; [226] developed the penalized dynamic programming algorithm to detect coefficient changes in discrete-time high-dimensional self-exciting Poisson processes in an offline setting.

This topic is also related to the multisource quickest detection problem, mostly assuming independence between multiple data streams. For instance, the quickest detection of the minimum of change-points for two independent compound Poisson processes was considered in [21] and optimal Bayesian sequential detection procedures were developed.

I. Change Detection-Isolation-Identification

In addition to detecting the change quickly after it occurs, sometimes we are also interested in identifying the post-change model and/or isolating a subset of nodes within a large network affected by the change. In [47], an asymptotically optimal Bayesian detection-isolation scheme was proposed assuming the post-change model is one of the finitely many distinct alternatives. In a series of works, Nikiforov introduced a minimax optimal detection-isolation algorithm for stochastic dynamical systems [137], developed a recursive variant of the algorithm that achieves better computational efficiency [138], and provided an asymptotic lower bound for the mean detection-isolation delay with constraints on the probability of false isolation and the average time before a false alarm [139]. Natural generalizations of CUSUM and SR procedures for detection-isolation problems were discussed in [198]. See [194], [196] for more detailed overviews.

J. Alternative Performance Metrics

Other than what have been presented in this survey, many alternative performance metrics have also been considered. For instance, [161] investigated an exponential penalty of delay rather than a linear penalty (as used in the definition of CADD, for instance). Such performance measures can be more accurate, sometimes for financial applications. In these cases, the change-point may not represent a time at which a fundamental shift in the performance occurs, but the compounding of investment growth can be a more suitable measure of the cost of delay. Similarly, in the health monitoring of components in aircraft systems, communication networks, and power grids, the effects of undetected faults can exponentiate with time. For problems involving estimation, the performance measures can also involve estimation accuracy, for instance, change-point location and other parameters involved in the problem. With many parallel data streams, the error metric can be the false discovery rate (FDR), which is the expected ratio of the number of falsely declared data streams to the total number of declared data streams [33].

IV. NEW DIMENSIONS

A. Machine Learning and Change Detection

Modern machine learning approaches can be adopted for solving sequential change detection problems, which we will review in this subsection.

1) *Density Ratio Estimation*: Instead of estimating the post-change density f_1 as in the GLR procedure, we may estimate the density ratio f_1/f_0 directly (referred to as density ratio estimation [192]), based on which we develop sequential change detection procedures. A data-driven framework using neural networks was developed in [134]. More specifically, given two sets of data sampled from the densities of interest, an optimization problem is defined so that the solution, specified through neural networks, will correspond to the desired likelihood ratio function or its transformations and can then be used for sequential change detection.

2) *Anomaly Detection*: Change detection is closely related to *anomaly detection*, which is a popular topic in machine learning and data mining, and many machine learning techniques have been developed. In particular, an recurrent neural network (RNN) based approach computes the detection statistic (referred to as the anomaly score) in an online fashion and compares with a threshold for anomaly detection [177]. The RNN-based approach can benefit certain situations since they are known to capture complex temporal dependencies for multivariate time series. We refer to [30] for a recent survey on deep learning techniques for anomaly detection. Developing mathematical theory for RNN-based sequential change detection is still an open question.

3) *Online Learning and Change Detection*: Online implementation is one of the most critical aspects of sequential change detection algorithms in practice. Although many algorithms enjoy recursive structure, such as CUSUM and SR procedures, some sequential detection procedures face a significant hurdle of online implementation due to their non-recursive nature. For instance, window-limited GLR statistic,

although enjoying robust performance in the presence of unknown post-change distributions, is not recursive since the parameters need to be continuously estimated by incorporating new samples. To tackle this challenge, inspired by online learning, [26] develops an online mirror descent-based GLR procedure to update the estimate of the unknown post-change parameter with new data. Another highly cited work [2] develops an online change detection procedure based on Bayesian computing. In recent work, [208] develops a framework for joint sequential change detection and online model fitting, which will be particularly suitable for parameterized models. A GLR procedure is developed in this framework using estimates of the unknown high-dimensional parameter obtained by the gradient descent update.

4) *Tracking Data Dynamics*: Many sequential data are dynamic even before the change has happened; for instance, solar flare detection from satellite video streaming [233], [234]. To build methods that work with real-world scenarios, we need to develop robust methods that can adapt to normal data dynamics without mislabeling them as change-points. A possible strategy is to combine tracking with detection. For instance, [233], [234] developed a procedure to detect sparse changes when the pre-change high-dimensional data is time-varying. The data dynamic is captured by tracking a time-varying manifold using variants of subspace tracking (e.g., GROUSE [245], PETRELS [39], or MOUSSE algorithm [234]). Another instance is the network Hawkes process model, where we may track the Hawkes process through online learning techniques [67].

5) *Active Learning and Change Detection*: For certain applications such as material science and recovering seafloor depth, data acquisition is expensive. Thus, it is desirable to collect data that is most useful in a sequential fashion, which is the theme of active learning (see, e.g., [29], [190]). The combination of active learning and change detection was introduced as active change-point detection (ACPD) problem in [74]. The task is to adaptively determine the next input to detect the change-point in a black-box expensive-to-evaluate function, with as few evaluations as possible. The method utilizes the existing change detection method to compute change scores and a Bayesian optimization method to determine the next input. A CUSUM procedure with an adaptive sampling strategy to detect mean shifts was developed in [120].

6) *Detection with Data Privacy*: As data privacy has growing importance in modern applications in social settings, it also leads to developing private change detection algorithms. Both offline and online change detection methods through the lens of *differential privacy* have been developed in [44]. A different privacy-aware sequential change detection method was studied in [104], using *maximal leakage* as the privacy metric, which is a weaker form of privacy compared with [44].

7) *Change Detection for Reinforcement Learning*: Reinforcement learning is a major type of sequential decision-making methodology in the era of artificial intelligence. How to implement reinforcement learning in a non-stationary and changing environment is still a mostly unexplored area. Recently, there have been some attempts to combine sequential change detection and reinforcement learning [147], where

change detection algorithms are utilized to detect the transition of the environment and trigger transitions of reinforcement learning algorithms.

B. Distribution-free Methods

Distribution-free methods aim to detect the change without making explicit distributional assumptions on the data. Such methods are particularly attractive in machine learning, such as kernel MMD based method discussed in Section III-B2, due to their flexibility in working with complex data. There have been kernel-based non-parametric methods developed in terms of change detection, both for the offline setting [8], [70], [72] and the online setting [115]. MMD statistics have also been used for anomalous sequence detection, for instance, [252], [254]. Besides MMD, other distribution-free methods have been developed for change detection. For instance, dissimilarity measures based on the kernel support vector machine (SVM) were built in [50], and generalized likelihood test directly using data empirical distributions when the true distributions are supported on a finite alphabet were constructed in [24], [105], [140], [141].

There are many other types of distribution-free non-parametric tests for change detection developed in various contexts. For instance, the maximal k -largest sample coherence between columns of each observed random matrix was developed to detect change for large-scale random matrices [12]. A nearest-neighbors-based statistic was proposed in [32] to detect the change in sequences of multivariate observations or non-Euclidean data objects such as network data. The weighted moving averages were studied in [58] to detect univariate drifts. A non-parametric approach was developed in [150] to detect departure from the reference signal with non-i.i.d. underlying time series. The spectral scan statistic for change detection over graphs was considered in [178]. Wasserstein distance was used to detect segments of times series in [38]. In [10], test statistics were constructed using martingales under the null hypothesis, and the rejection threshold is determined using a uniform non-asymptotic law of the iterated logarithm.

C. Non-stationary Multi-armed Bandits with Changes

Multi-armed bandit is a class of fundamental problems in online learning and sequential decision-making. A learning agent aims to maximize its expected cumulative reward by repeatedly selecting to pull one arm at each time step. Change detection can play a role in the scenario where the reward distributions may change in a piece-wise-stationary fashion at unknown time steps. To handle dynamic multi-armed bandit problems, various change detection methods were considered, including the Page-Hinkley test [73], a windowed mean-shift detection [243], CUSUM test [119], and sample mean based test [25]. Usually, the algorithm will reset once a change is detected. From a Bayesian perspective, the Thompson sampling strategy equipped with a Bayesian change-point mechanism was considered in [128]. The adversarial multi-armed bandit problem with change points was also considered in [5].

D. Optimization for Change Detection and Estimation

Optimization is becoming a centerpiece in developing modern machine learning algorithms. Recent advances in convex optimization have enabled solving many large-scale problems. A line of research aims to cast (offline) change detection and estimation (of their locations) as an optimization problem. The benefits of this optimization-based approach typically include computational efficiency (when the optimization problem is convex) and theoretical performance guarantees based on optimization theory. Below we give some examples.

The univariate change detection for a mean shift using an optimization approach has been studied in [117], and performance guarantees were established by relating the ℓ_2 recovery error to detection performance. A ℓ_0 -penalized least squares method was considered in [224]. By connecting to binary segmentation methods, change detection and localization for univariate data in the non-parametric settings was studied in [148].

Multivariate change detection using an optimization approach has also been studied. For instance, a dynamic programming approach was developed for recovering an unknown number of change-points from multivariate autoregressive models [225]. A network binary segmentation method for change detection was proposed in [223], which has been extended for covariance matrix change detection in [222]. Finally, the work [191] combined the filtered derivative with convex optimization methods to estimate change-points for multi-dimensional data.

V. MODERN APPLICATIONS

Sequential change detection has traditionally been used in industrial process monitoring applications, which was probably the original motivation for change detection procedures to be developed in the early days. The wide adoption of change detection in industrial quality engineering and manufacturing initiates the field of statistical process control (SPC) (see, e.g., [143], [182]). Recently, there have been many more modern applications for sequential change detection, and we present a selection of them here.

A. Smart Grids

The sequential change detection methodology has been recently successfully applied for sequential line outage detection in power transmission systems. In modern smart grids, high-speed synchronized voltage phase angle measurements are taken from phasor measurement units (PMU). Based on PMU measurements, a linearized incremental small-signal power system model was developed in [37]. Once a line outage occurs, there is a change in the covariance matrix of incremental phases, by monitoring which, line outages can be detected and identified using sequential change detection algorithms. In [171], the transient dynamics of the power system following a line outage is further incorporated. The D-CUSUM algorithm was then developed to incorporate the dynamic nature of the line outage in [171] (see Section III-C1 for more details).

There have been other works on sequential change detection for smart grids. The generalized local likelihood ratio test was applied for voltage quality monitoring [114], photovoltaic systems [36], attack detection in the multi-agent reputation systems [113], wide-area monitoring [112], and cyber-attacks detection in discrete-time linear dynamic system [92], [93]. The decentralized detection with level-triggered sampling was considered in [241]. In [77], a general stochastic graphical framework for modeling the bus measurements and a data-adaptive data-acquisition and decision-making processes were designed for the quickest search and localization of anomaly in power grids.

B. Cybersecurity

Cybersecurity has become a critical problem with the development of wireless communication, networking, and the Internet of things. It is of practical importance to detect attacks and intrusions in real-time from network streaming data, e.g., denial-of-service attacks, worm-based attacks, port-scanning, and man-in-the-middle attacks. The sequential change detection approach is a natural fit since the attacks usually change network traffic distribution. In [203], multi-channel generalizations of the CUSUM procedure and non-parametric tests were proposed. In [204], adaptive sequential methods were proposed for early detection of subtle network attacks, utilizing data from multiple layers of the network protocol. In [202], a multi-cyclic detection procedure based on the SR procedure was proposed. In [199], score-based CUSUM and SR procedures were exploited for network anomaly detection, and a hybrid detection system was proposed. The application to cybersecurity was also discussed in books [194], [196], and recent reviews [84], [85].

C. Sensors Networks

Sensor networks collecting sequential data have been widely used for geophysical, environmental, traffic, and internet traffic monitoring applications, which we will briefly summarize in this subsection.

Seismology is experiencing rapid growth in the quantity of data. Earthquake detection aims to identify seismic events in continuous data – a fundamental operation for seismology [242]. Modern ultra-dense seismic sensor arrays have obtained a massive amount of continuous data for seismic studies, and many such data are publicly available through IRIS [1]. In the old days, network seismology treated seismic signals individually - one sensor at a time - and detected an earthquake upon multiple impulsive arrivals consistent with a source within the Earth [87]. Recently, with advances in sensor technology, which bring densely sampled data, high-performance computing and high-speed communication, we are able to use a *network-based detection* by exploiting correlations between sensors to extract coherence signals. This will enhance the systematic detection of weak and unusual events that currently go undetected using individual sensors. Detecting such weak events is very crucial for earthquake prediction [145], [219], oil field exploration, volcano monitoring, and deeper earth studies [80]. Towards this goal, in [232], a subspace-CUSUM

procedure was developed for network-based detection by exploiting the low-rank subspace structure induced by waveform similarity.

Sensor networks have also been deployed to monitor drinking water safety from the water tower to private residences. Sequential change detection using residual chlorine concentration measurements from the sensors network was developed in [65]. Methods have also been developed for monitoring river contamination [34], [35], which specifically consider the spatio-temporal correlation in observations along the sensor network due to water dynamics.

Sequential monitoring of traffic flow using traffic sensors has been considered in [166], and a distributed, online, sequential algorithm for detecting multiple faults in a sensor network was presented therein. Recently, Hawkes processes models for correlated traffic anomalies using data collected by inductive-loop traffic detectors were developed in [250].

D. Wireless Communications

Sequential change detection has been used for wireless communications, including online user activity detection for multi-user direct-sequence/code-division multiple-access (DS-CDMA) environment [146], detecting “spectrum opportunities” in the cognitive radio setting by identifying the occupancy and idle of channels from primary user’s activities [95], [235]. More recently, [81] established a change detection framework for low probability of detection (LPD) communication, where a transmitter, Alice, wants to hide her transmission to a receiver, Bob, from an adversary, Willie; three different sequential tests were considered, including Shewhart, CUSUM, and SR procedures, to model Willie’s detection process.

E. Video Processing and Computer Vision

Change detection is one of the most commonly encountered low-level tasks in computer vision and video processing [163], and many such problems are essentially sequential. A plethora of practical algorithms have been developed to date; for instance, scene change detection [110], street-view change detection [3], and change detection in video sequences [211]. In [48], a pixel-based weightless neural network (WNN) method was developed to detect changes in the field of view of a camera. In [118], multiple images from reference and mission passes of a scene of interest were used to improve detection performance. There are still many open questions regarding how to leverage the power of statistical sequential change detection for computer vision and video processing. We present an example of solar flare detection from video sequences in Fig. 4, which has been considered in several works along this line including [234].

F. Social Networks

The wide-spread use of social networks and the great availability of information networks (e.g., Twitter, Facebook, blogs) lead to a large amount of user-generated data [91], which are quite valuable in studying many social phenomena. One important aspect is to detect change-points in streaming

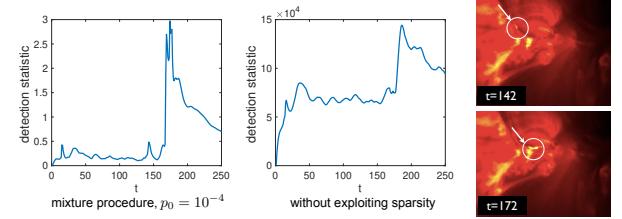


Fig. 4. Solar flare detection with the mixture procedure as considered in [234]; the first minor solar flare at $t = 142$ is hardly visible, and it is missed entirely by a baseline detection statistic (sum of CUSUM at each data dimension without exploiting sparsity in the change: “solar flare”). This also illustrates the importance of exploiting sparsity in the change.

social network data [90], which may represent the collective anticipation of or response to external events or system “shocks” [152]. Detecting such changes could provide a better understanding of the patterns of social life. In other cases, early detection of change-points can predict or even prevent social stress due to disease or international threat, for instance, detecting self-exciting changes (modeled by network Hawkes processes) in social networks [116]. A related topic is distributed hypothesis testing in social networks: [101] showed the exponential convergence rate of a Bayesian update scheme of nodal belief (distribution estimate) in the social learning setting.

G. Epidemiology

Sequential change detection can potentially play an important role in public health and disease surveillance. Early detection of epidemics is a very important topic. In [17], [18], Baron cast the early detection of epidemics as a Bayes sequential change detection problem and proposed an asymptotically pointwise optimal stopping rule, which is computationally efficient for complicated prior distributions arising in epidemiology. In [244], a modified CUSUM procedure was proposed for the susceptible–infected–recovered (SIR) epidemic model to detect change-point in the infection rate parameter. Moreover, change detection has been incorporated into studying the intervention’s effectiveness, based on the premise that the underlying epidemiological model may change over time due to interventions. Evaluating intervention measures’ effectiveness requires detecting underlying change-points, which becomes even more important in the COVID-19 era [249]. Such works include [151], [228], which estimate the change-points in time series to assess the effectiveness of interventions such as lock-down and mask usage; in [49], the problem of detecting the growth rate change for the COVID-19 spread in Germany was studied, where results were further incorporated into forecasting. There are still many open questions in this area regarding developing effective sequential change detection procedures suitable for infectious disease early detection.

VI. CONCLUSIONS

Our goal in this survey was to provide a glimpse of the past and recent advances in sequential change detection, and its application in various domains. We have covered different types

of sequential change detection procedures, both theoretically optimal and practical. We also discussed how the intersection of sequential change detection with other areas has created interesting new directions for research.

ACKNOWLEDGEMENT

The authors are grateful to the Guest Editor and the anonymous reviewers for their helpful comments.

REFERENCES

- [1] “Incorporated research institutions for seismology,” <https://www.iris.edu>, accessed: 2020-10-12.
- [2] R. P. Adams and D. J. MacKay, “Bayesian online changepoint detection,” *arXiv:0710.3742*, 2007.
- [3] P. F. Alcantarilla, S. Stent, G. Ros, R. Arroyo, and R. Gherardi, “Street-view change detection with deconvolutional networks,” *Autonomous Robots*, vol. 42, no. 7, pp. 1301–1322, 2018.
- [4] D. Aldous, *Probability Approximations via the Poisson Clumping Heuristic*. Springer Science & Business Media, 2013, vol. 77.
- [5] R. Allesiardo and R. Féraud, “Exp3 with drift detection for the switching bandit problem,” in *Proceedings of the IEEE International Conference on Data Science and Advanced Analytics (DSAA)*, 2015, pp. 1–7.
- [6] S. Aminikhahhahi and D. J. Cook, “A survey of methods for time series change point detection,” *Knowledge and Information Systems*, vol. 51, no. 2, pp. 339–367, 2017.
- [7] D. Amorese, “Applying a change-point detection method on frequency-magnitude distributions,” *Bulletin of the Seismological Society of America*, vol. 97, no. 5, pp. 1742–1749, 2007.
- [8] S. Arlot, A. Celisse, and Z. Harchaoui, “A kernel multiple change-point algorithm via model selection.” *Journal of Machine Learning Research*, vol. 20, no. 162, pp. 1–56, 2019.
- [9] J. Bai, “Estimating multiple breaks one at a time,” *Econometric Theory*, vol. 13, no. 3, pp. 315–352, 1997.
- [10] A. Balsubramani and A. Ramdas, “Sequential nonparametric testing with the law of the iterated logarithm,” in *Proceedings of the Conference on Uncertainty in Artificial Intelligence (UAI)*, 2016, pp. 42–51.
- [11] T. Banerjee and V. V. Veeravalli, “Data-efficient quickest change detection with on-off observation control,” *Sequential Analysis*, vol. 31, no. 1, pp. 40–77, 2012.
- [12] T. Banerjee, H. Firouzi, and A. O. Hero, “Quickest detection for changes in maximal kNN coherence of random matrices,” *IEEE Transactions on Signal Processing*, vol. 66, no. 17, pp. 4490–4503, 2018.
- [13] T. Banerjee and V. V. Veeravalli, “Data-efficient quickest change detection in minimax settings,” *IEEE Transactions on Information Theory*, vol. 59, no. 10, pp. 6917–6931, 2013.
- [14] —, “Data-efficient minimax quickest change detection with composite post-change distribution,” *IEEE Transactions on Information Theory*, vol. 61, no. 9, pp. 5172–5184, 2015.
- [15] —, “Data-efficient quickest change detection in sensor networks,” *IEEE Transactions on Signal Processing*, vol. 63, no. 14, pp. 3727–3735, 2015.
- [16] I. Barnett and J.-P. Onnela, “Change point detection in correlation networks,” *Scientific Reports*, vol. 6, 2015.
- [17] M. Baron, “Early detection of epidemics as a sequential change-point problem,” in *Longevity, Aging, and Degradation Models in Reliability, Public Health, Medicine and Biology*, 2004, pp. 31–43.
- [18] —, “Bayes and asymptotically pointwise optimal stopping rules for the detection of influenza epidemics,” in *Case Studies in Bayesian Statistics*. Springer, 2002, pp. 153–163.
- [19] M. Basseville and I. V. Nikiforov, *Detection of Abrupt Changes: Theory and Application*. Prentice Hall, 1993.
- [20] E. Bayraktar and L. Lai, “Byzantine fault tolerant distributed quickest change detection,” *SIAM Journal on Control and Optimization*, vol. 53, no. 2, pp. 575–591, 2015.
- [21] E. Bayraktar and H. V. Poor, “Quickest detection of a minimum of two Poisson disorder times,” *SIAM Journal on Control and Optimization*, vol. 46, no. 1, pp. 308–331, 2007.
- [22] I. Berkes, E. Gombay, L. Horváth, and P. Kokoszka, “Sequential change-point detection in GARCH(p, q) models,” *Econometric Theory*, vol. 20, no. 6, pp. 1140–1167, 2004.
- [23] L. Boysen, A. Kempe, V. Liebscher, A. Munk, and O. Wittich, “Consistencies and rates of convergence of jump-penalized least squares estimators,” *Annals of Statistics*, vol. 37, no. 1, pp. 157–183, 2009.
- [24] Y. Bu, S. Zou, and V. V. Veeravalli, “Linear-complexity exponentially-consistent tests for universal outlying sequence detection,” *IEEE Transactions on Signal Processing*, vol. 67, no. 8, pp. 2115–2128, 2019.
- [25] Y. Cao, Z. Wen, B. Kveton, and Y. Xie, “Nearly optimal adaptive procedure with change detection for piecewise-stationary bandit,” in *Proceedings of the International Conference on Artificial Intelligence and Statistics (AISTATS)*, 2019, pp. 418–427.
- [26] Y. Cao, L. Xie, Y. Xie, and H. Xu, “Sequential change-point detection via online convex optimization,” *Entropy*, vol. 20, no. 2, 2018.
- [27] Y. Cao and Y. Xie, “Robust sequential change-point detection by convex optimization,” in *Proceedings of the IEEE International Symposium on Information Theory (ISIT)*, 2017, pp. 1287–1291.
- [28] Y. Cao, Y. Xie, and N. Gebraeel, “Multi-sensor slope change detection,” *Annals of Operations Research*, vol. 263, no. 1-2, pp. 163–189, 2018.
- [29] R. Castro and R. Nowak, “Active learning and sampling,” in *Foundations and Applications of Sensor Management*. Springer, 2008, pp. 177–200.
- [30] R. Chalapathy and S. Chawla, “Deep learning for anomaly detection: A survey,” *arXiv:1901.03407*, 2019.
- [31] J.-F. Chamberland and V. V. Veeravalli, “Wireless sensors in distributed detection applications,” *IEEE Signal Processing Magazine Special Issue on Resource-Constrained Signal Processing, Communications, and Networking*, vol. 24, no. 3, pp. 16–25, May 2007.
- [32] H. Chen, “Sequential change-point detection based on nearest neighbors,” *Annals of Statistics*, vol. 47, no. 3, pp. 1381–1407, 2019.
- [33] J. Chen, W. Zhang, and H. V. Poor, “A false discovery rate oriented approach to parallel sequential change detection problems.” *IEEE Transactions on Signal Processing*, vol. 68, pp. 1823–1836, 2020.
- [34] J. Chen, S.-H. Kim, and Y. Xie, “S³T: An efficient score-statistic for spatio-temporal surveillance,” *arXiv:1706.05331*, 2018.
- [35] J. Chen, C. Park, S.-H. Kim, and Y. Xie, “To reduce or not to reduce: A study on spatio-temporal surveillance,” *Environmental and Ecological Statistics*, vol. 26, no. 3, pp. 217–238, 2019.
- [36] L. Chen, S. Li, and X. Wang, “Quickest fault detection in photovoltaic systems,” *IEEE Transactions on Smart Grid*, vol. 9, no. 3, pp. 1835–1847, 2016.
- [37] Y. C. Chen, T. Banerjee, A. D. Dominguez-Garcia, and V. V. Veeravalli, “Quickest line outage detection and identification,” *IEEE Transactions on Power Systems*, vol. 31, no. 1, pp. 749–758, 2015.
- [38] K. C. Cheng, S. Aeron, M. C. Hughes, E. Hussey, and E. L. Miller, “Optimal transport based change point detection and time series segment clustering,” in *Proceedings of the IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2020, pp. 6034–6038.
- [39] Y. Chi, Y. C. Eldar, and R. Calderbank, “PETRELS: Parallel subspace estimation and tracking using recursive least squares from partial observations,” *IEEE Transactions on Signal Processing*, vol. 61, no. 23, pp. 5947–5959, 2013.
- [40] H. Cho and P. Fryzlewicz, “Multiscale and multilevel technique for consistent segmentation of nonstationary time series,” *Statistica Sinica*, vol. 22, no. 1, pp. 207–229, 2012.
- [41] —, “Multiple-change-point detection for high dimensional time series via sparsified binary segmentation,” *Journal of the Royal Statistical Society. Series B (Statistical Methodology)*, vol. 77, no. 2, pp. 475–507, 2015.
- [42] H. Choi, H. Ombao, and B. Ray, “Sequential change-point detection methods for nonstationary time series,” *Technometrics*, vol. 50, no. 1, pp. 40–52, 2008.
- [43] Y. S. Chow, H. Robbins, and D. Siegmund, *Great Expectations: The Theory of Optimal Stopping*. Houghton Mifflin, 1971.
- [44] R. Cummings, S. Krehbiel, Y. Mei, R. Tuo, and W. Zhang, “Differentially private change-point detection,” in *Proceedings of the Advances in Neural Information Processing Systems (NeurIPS)*, 2018, pp. 10 825–10 834.
- [45] D. J. Daley and D. Vere-Jones, *An Introduction to The Theory of Point Processes: Volume II: General Theory and Structure*. Springer Science & Business Media, 2007.
- [46] S. Dayanik, C. Goulding, and H. V. Poor, “Bayesian sequential change diagnosis,” *Mathematics of Operations Research*, vol. 33, no. 2, pp. 475–496, 2008.
- [47] S. Dayanik, W. B. Powell, and K. Yamazaki, “Asymptotically optimal Bayesian sequential change detection and identification rules,” *Annals of Operations Research*, vol. 208, no. 1, pp. 337–370, 2013.

- [48] M. De Gregorio and M. Giordano, "Change detection with weightless neural networks," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition Workshop*, 2014, pp. 403–407.
- [49] J. Dehning, J. Zierenberg, F. P. Spitzner, M. Wibral, J. P. Neto, M. Wilczek, and V. Priesemann, "Inferring change points in the spread of COVID-19 reveals the effectiveness of interventions," *Science*, vol. 369, no. 6500, 2020.
- [50] F. Desobry, M. Davy, and C. Doncarli, "An online kernel change detection algorithm," *IEEE Transactions on Signal Processing*, vol. 53, no. 8, pp. 2961–2974, 2005.
- [51] E. Ebrahimbazeh and A. Tchamkerten, "Sequential detection of transient changes in stochastic systems under a sampling constraint," in *Proceedings of the IEEE International Symposium on Information Theory (ISIT)*, 2015, pp. 156–160.
- [52] D. Egea-Roca, J. A. López-Salcedo, G. Seco-Granados, and H. V. Poor, "Performance bounds for finite moving average tests in transient change detection," *IEEE Transactions on Signal Processing*, vol. 66, no. 6, pp. 1594–1606, 2018.
- [53] F. Enikeeva and Z. Harchaoui, "High-dimensional change-point detection under sparse alternatives," *Annals of Statistics*, vol. 47, no. 4, pp. 2051–2079, 2019.
- [54] M. Farajtabar, Y. Wang, M. Gomez-Rodriguez, S. Li, H. Zha, and L. Song, "COEVOLVE: A joint point process model for information diffusion and network evolution," *Journal of Machine Learning Research*, vol. 18, no. 1, pp. 1305–1353, 2017.
- [55] G. Fellouris and A. Tartakovsky, "Multichannel sequential detection—Part I: Non-iid data," *IEEE Transactions on Information Theory*, vol. 63, no. 7, pp. 4551–4571, 2017.
- [56] G. Fellouris, E. Bayraktar, and L. Lai, "Efficient Byzantine sequential change detection," *IEEE Transactions on Information Theory*, vol. 64, no. 5, pp. 3346–3360, 2017.
- [57] E. W. Fox, M. B. Short, F. P. Schoenberg, K. D. Coronges, and A. L. Bertozzi, "Modeling e-mail networks and inferring leadership using self-exciting point processes," *Journal of the American Statistical Association*, vol. 111, no. 514, pp. 564–584, 2016.
- [58] I. Frías-Blanco, J. del Campo-Avila, G. Ramos-Jimenez, R. Morales-Bueno, A. Ortiz-Díaz, and Y. Caballero-Mota, "Online and non-parametric drift detection methods based on Hoeffding's bounds," *IEEE Transactions on Knowledge and Data Engineering*, vol. 27, no. 3, pp. 810–823, 2014.
- [59] K. Frick, A. Munk, and H. Sieling, "Multiscale change point inference," *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, vol. 76, no. 3, pp. 495–580, 2014.
- [60] P. Fryzlewicz and S. S. Rao, "Multiple-change-point detection for autoregressive conditional heteroscedastic processes," *Journal of the Royal Statistical Society. Series B (Statistical Methodology)*, vol. 76, no. 5, pp. 903–924, 2014.
- [61] P. Fryzlewicz, "Wild binary segmentation for multiple change-point detection," *Annals of Statistics*, vol. 42, no. 6, pp. 2243–2281, 2014.
- [62] C. Fuh and A. G. Tartakovsky, "Asymptotic Bayesian theory of quickest change detection for hidden Markov models," *IEEE Transactions on Information Theory*, vol. 65, no. 1, pp. 511–529, 2019.
- [63] J. Geng, B. Zhang, L. M. Huie, and L. Lai, "Online change-point detection of linear regression models," *IEEE Transactions on Signal Processing*, vol. 67, no. 12, pp. 3316–3329, 2019.
- [64] B. K. Guépié, L. Fillatre, and I. Nikiforov, "Sequential detection of transient changes," *Sequential Analysis*, vol. 31, no. 4, pp. 528–547, 2012.
- [65] B. K. Guepie, L. Fillatre, and I. Nikiforov, "Sequential monitoring of water distribution network," *IFAC Proceedings Volumes*, vol. 45, no. 16, pp. 392–397, 2012.
- [66] O. Hadjiliadis, H. Zhang, and H. V. Poor, "One shot schemes for decentralized quickest change detection," *IEEE Transactions on Information Theory*, vol. 55, no. 7, pp. 3346–3359, 2009.
- [67] E. C. Hall and R. M. Willett, "Tracking dynamic point processes on networks," *IEEE Transactions on Information Theory*, vol. 62, no. 7, pp. 4327–4346, 2016.
- [68] D. Hallac, J. Leskovec, and S. Boyd, "Network lasso: Clustering and optimization in large graphs," in *Proceedings of the International Conference on Knowledge Discovery and Data Mining*, 2015, pp. 387–396.
- [69] Z. Harchaoui and O. Cappe, "Retrospective multiple change-point estimation with kernels," in *IEEE/SP 14th Workshop on Statistical Signal Processing*, 2007, pp. 768–772.
- [70] Z. Harchaoui, F. Bach, O. Cappe, and E. Moulines, "Kernel-based methods for hypothesis testing: A unified view," *IEEE Signal Processing Magazine*, vol. 30, no. 4, pp. 87–97, 2013.
- [71] Z. Harchaoui and C. Lévy-Leduc, "Multiple change-point estimation with a total variation penalty," *Journal of the American Statistical Association*, vol. 105, no. 492, pp. 1480–1493, 2010.
- [72] Z. Harchaoui, E. Moulines, and F. R. Bach, "Kernel change-point analysis," in *Proceedings of the Advances in Neural Information Processing Systems (NeurIPS)*, 2009, pp. 609–616.
- [73] C. Hartland, N. Baskiotis, S. Gelly, M. Sebag, and O. Teytaud, "Change point detection and meta-bandits for online learning in dynamic environments," in *CAp: 9è Conférence francophone sur l'apprentissage automatique*, 2007, pp. 237–250.
- [74] S. Hayashi, Y. Kawahara, and H. Kashima, "Active change-point detection," in *Asian Conference on Machine Learning*, 2019, pp. 1017–1032.
- [75] X. He, T. Rekatsinas, J. Foulds, L. Getoor, and Y. Liu, "HawkesTopic: A joint model for network inference and topic modeling from text-based cascades," in *Proceedings of the International Conference on Machine Learning (ICML)*, 2015.
- [76] T. Herberth and U. Jensen, "Optimal detection of a change point in a Poisson process for different observation schemes," *Scandinavian Journal of Statistics*, vol. 31, no. 3, pp. 347–366, 2004.
- [77] J. Heydari and A. Tajer, "Quickest localization of anomalies in power grids: A stochastic graphical framework," *IEEE Transactions on Smart Grid*, vol. 9, no. 5, pp. 4679–4688, 2017.
- [78] ———, "Quickest search and learning over correlated sequences: Theory and application," *IEEE Transactions on Signal Processing*, vol. 67, no. 3, pp. 638–651, 2018.
- [79] ———, "Quickest search for a change point," in *Proceedings of the IEEE International Symposium on Information Theory (ISIT)*, 2019, pp. 2194–2198.
- [80] H. H. Huang, F.-C. Lin, B. Schmandt, J. Farrell, R. B. Smith, and V. C. Tsai, "The Yellowstone magmatic system from the mantle plume to the upper crust," *Science*, vol. 348, pp. 773–776, 2015.
- [81] K.-W. Huang, H.-M. Wang, D. Towsley, and H. V. Poor, "LPD communication: A sequential change-point detection perspective," *IEEE Transactions on Communications*, vol. 68, no. 4, pp. 2474–2490, 2020.
- [82] P. J. Huber, "A robust version of the probability ratio test," *Annals of Mathematical Statistics*, vol. 36, no. 6, pp. 1753–1758, 1965.
- [83] ———, *Robust Statistics*. New York, NY, USA: Wiley, 1981.
- [84] D. R. Jeske, N. T. Stevens, A. G. Tartakovsky, and J. D. Wilson, "Statistical methods for network surveillance," *Applied Stochastic Models in Business and Industry*, vol. 34, no. 4, pp. 425–445, 2018.
- [85] D. R. Jeske, N. T. Stevens, J. D. Wilson, and A. G. Tartakovsky, "Statistical network surveillance," *Wiley StatsRef: Statistics Reference Online*, pp. 1–12, 2014.
- [86] Y. Jiao, Y. Chen, and Y. Gu, "Subspace change-point detection: A new model and solution," *IEEE Journal of Selected Topics in Signal Processing*, vol. 12, no. 6, pp. 1224–1239, 2018.
- [87] C. E. Johnson and A. G. L. Hirshorn, "Robust regional phase associations," USGS Open File Report 94-621, Tech. Rep., 1994.
- [88] I. M. Johnstone, "On the distribution of the largest eigenvalue in principal components analysis," *Annals of statistics*, vol. 29, no. 2, pp. 295–327, 2001.
- [89] R. Killick, P. Fearnhead, and I. A. Eckley, "Optimal detection of changepoints with a linear computational cost," *Journal of the American Statistical Association*, vol. 107, no. 500, pp. 1590–1598, 2012.
- [90] V. Krishnamurthy, "Quickest detection POMDPs with social learning: Interaction of local and global decision makers," *IEEE Transactions on Information Theory*, vol. 58, no. 8, pp. 5563–5587, 2012.
- [91] V. Krishnamurthy and H. V. Poor, "A tutorial on interactive sensing in social networks," *IEEE Transactions on Computational Social Systems*, vol. 1, no. 1, pp. 3–21, 2014.
- [92] M. N. Kurt, Y. Yilmaz, and X. Wang, "Distributed quickest detection of cyber-attacks in smart grid," *IEEE Transactions on Information Forensics and Security*, vol. 13, no. 8, pp. 2015–2030, 2018.
- [93] ———, "Real-time detection of hybrid and stealthy cyber-attacks in smart grid," *IEEE Transactions on Information Forensics and Security*, vol. 14, no. 2, pp. 498–513, 2019.
- [94] M. N. Kurt and X. Wang, "Multisensor sequential change detection with unknown change propagation pattern," *IEEE Transactions on Aerospace and Electronic Systems*, vol. 55, no. 3, pp. 1498–1518, 2018.
- [95] L. Lai, Y. Fan, and H. V. Poor, "Quickest detection in cognitive radio: A sequential change detection framework," in *Proceedings of the IEEE Global Telecommunications Conference*, 2008, pp. 1–5.
- [96] T. L. Lai, "Information bounds and quick detection of parameter changes in stochastic systems," *IEEE Transactions on Information Theory*, vol. 44, no. 7, pp. 2917–2929, 1998.

- [97] ——, “Sequential changepoint detection in quality control and dynamical systems,” *Journal of the Royal Statistical Society: Series B (Methodological)*, vol. 57, no. 4, pp. 613–644, 1995.
- [98] ——, “Sequential analysis: Some classical problems and new challenges,” *Statistica Sinica*, vol. 11, no. 2, pp. 303–351, 2001.
- [99] T. L. Lai and J. Z. Shan, “Efficient recursive algorithms for detection of abrupt changes in signals and control systems,” *IEEE Transactions on Automatic Control*, vol. 44, no. 5, pp. 952–966, 1999.
- [100] A. Lakhina, M. Crovella, and C. Diot, “Diagnosing network-wide traffic anomalies,” *ACM SIGCOMM Computer Communication Review*, vol. 34, no. 4, pp. 219–230, 2004.
- [101] A. Lalitha, T. Javidi, and A. D. Sarwate, “Social learning and distributed hypothesis testing,” *IEEE Transactions on Information Theory*, vol. 64, no. 9, pp. 6161–6179, 2018.
- [102] L. Lamport, R. Shostak, and M. Pease, “The Byzantine generals problem,” in *Concurrency: the Works of Leslie Lamport*, 2019, pp. 203–226.
- [103] T. S. Lau and W. P. Tay, “Quickest change detection in the presence of a nuisance change,” *IEEE Transactions on Signal Processing*, vol. 67, no. 20, pp. 5281–5296, 2019.
- [104] ——, “Privacy-aware quickest change detection,” in *Proceedings of the IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2020, pp. 5999–6003.
- [105] T. S. Lau, W. P. Tay, and V. V. Veeravalli, “A binning approach to quickest change detection with unknown post-change distribution,” *IEEE Transactions on Signal Processing*, vol. 67, no. 3, pp. 609–621, 2018.
- [106] M. Lavielle, “Using penalized contrasts for the change-point problem,” *Signal Processing*, vol. 85, no. 8, pp. 1501–1510, 2005.
- [107] M. Lavielle and E. Moulines, “Least-squares estimation of an unknown number of shifts in a time series,” *Journal of Time Series Analysis*, vol. 21, no. 1, pp. 33–59, 2000.
- [108] É. Lebarbier, “Detecting multiple change-points in the mean of Gaussian process by model selection,” *Signal Processing*, vol. 85, no. 4, pp. 717–736, 2005.
- [109] K.-C. Lee and D. Kriegman, “Online learning of probabilistic appearance manifolds for video-based recognition and tracking,” in *Proceedings of the IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR)*, vol. 1, 2005, pp. 852–859.
- [110] D. Lelescu and D. Schonfeld, “Statistical sequential analysis for real-time video scene change detection on compressed multimedia bitstream,” *IEEE Transactions on Multimedia*, vol. 5, no. 1, pp. 106–117, 2003.
- [111] J. Li, D. Towsley, S. Zou, V. V. Veeravalli, and G. Ciocarlie, “A consensus-based approach for distributed quickest detection of significant events in networks,” in *Proceedings of the Asilomar Conference on Signals, Systems, and Computers*, 2019, pp. 1–4.
- [112] S. Li, Y. Yilmaz, and X. Wang, “Quickest detection of false data injection attack in wide-area smart grids,” *IEEE Transactions on Smart Grid*, vol. 6, no. 6, pp. 2725–2735, 2015.
- [113] S. Li and X. Wang, “Quickest attack detection in multi-agent reputation systems,” *IEEE Journal of Selected Topics in Signal Processing*, vol. 8, no. 4, pp. 653–666, 2014.
- [114] ——, “Cooperative change detection for voltage quality monitoring in smart grids,” *IEEE Transactions on Information Forensics and Security*, vol. 11, no. 1, pp. 86–99, 2015.
- [115] S. Li, Y. Xie, H. Dai, and L. Song, “Scan B-statistic for kernel changepoint detection,” *Sequential Analysis*, vol. 38, no. 4, pp. 503–544, 2019.
- [116] S. Li, Y. Xie, M. Farajtabar, A. Verma, and L. Song, “Detecting changes in dynamic events over networks,” *IEEE Transactions on Signal and Information Processing over Networks*, vol. 3, no. 2, pp. 346–359, 2017.
- [117] K. Lin, J. Sharpnack, A. Rinaldo, and R. J. Tibshirani, “Approximate recovery in changepoint problems, from ℓ_2 estimation error rates,” *arXiv:1606.06746*, 2016.
- [118] A. J. Lingg, E. Zelnio, F. Garber, and B. D. Rigling, “A sequential framework for image change detection,” *IEEE Transactions on Image Processing*, vol. 23, no. 5, pp. 2405–2413, 2014.
- [119] F. Liu, J. Lee, and N. Shroff, “A change-detection based framework for piecewise-stationary multi-armed bandit problem,” in *Proceedings of the AAAI Conference on Artificial Intelligence*, 2018, pp. 3651–3658.
- [120] K. Liu, Y. Mei, and J. Shi, “An adaptive sampling strategy for online high-dimensional process monitoring,” *Technometrics*, vol. 57, no. 3, pp. 305–319, 2015.
- [121] K. Liu, R. Zhang, and Y. Mei, “Scalable sum-shrinkage schemes for distributed monitoring large-scale data streams,” *Statistica Sinica*, vol. 29, no. 1, pp. 1–22, 2019.
- [122] G. Lorden, “Procedures for reacting to a change in distribution,” *Annals of Mathematical Statistics*, vol. 42, no. 6, pp. 1897–1908, 1971.
- [123] ——, “Open-ended tests for Koopman-Darmois families,” *Annals of Statistics*, vol. 1, no. 4, pp. 633–643, 1973.
- [124] M. Ludkovski, “Bayesian quickest detection with observation-changepoint feedback,” in *Proceedings of the IEEE Annual Conference on Decision and Control (CDC)*, 2012, pp. 166–171.
- [125] D. S. Matteson and N. A. James, “A nonparametric approach for multiple change point analysis of multivariate data,” *Journal of the American Statistical Association*, vol. 109, no. 505, pp. 334–345, 2014.
- [126] Y. Mei, “Efficient scalable schemes for monitoring a large number of data streams,” *Biometrika*, vol. 97, no. 2, pp. 419–433, 2010.
- [127] ——, “Information bounds and quickest change detection in decentralized decision systems,” *IEEE Transactions on Information theory*, vol. 51, no. 7, pp. 2669–2681, 2005.
- [128] J. Mellor and J. Shapiro, “Thompson sampling in switching environments with Bayesian online change detection,” in *Proceedings of the International Conference on Artificial Intelligence and Statistics (AISTATS)*, 2013, pp. 442–450.
- [129] T. L. Molloy and J. J. Ford, “Misspecified and asymptotically minimax robust quickest change detection,” *IEEE Transactions on Signal Processing*, vol. 65, no. 21, pp. 5730–5742, 2017.
- [130] D. C. Montgomery, *Introduction to Statistical Quality Control*. John Wiley & Sons, 2007.
- [131] P. Moulin and V. V. Veeravalli, *Statistical Inference for Engineers and Data Scientists*. Cambridge University Press, 2018.
- [132] G. V. Moustakides, “Optimal stopping times for detecting changes in distributions,” *Annals of Statistics*, vol. 14, no. 4, pp. 1379–1387, 1986.
- [133] ——, “Multiple optimality properties of the Shewhart test,” *Sequential Analysis*, vol. 33, no. 3, pp. 318–344, 2014.
- [134] G. V. Moustakides and K. Basioti, “Training neural networks for likelihood/density ratio estimation,” *arXiv:1911.00405*, 2019.
- [135] G. V. Moustakides, G. H. Jajamovich, A. Tajer, and X. Wang, “Joint detection and estimation: Optimum tests and applications,” *IEEE Transactions on Information Theory*, vol. 58, no. 7, pp. 4215–4229, 2012.
- [136] J. Neyman and E. S. Pearson, “IX. On the problem of the most efficient tests of statistical hypotheses,” *Philosophical Transactions of the Royal Society of London. Series A, Containing Papers of a Mathematical or Physical Character*, vol. 231, no. 694–706, pp. 289–337, 1933.
- [137] I. V. Nikiforov, “A generalized change detection problem,” *IEEE Transactions on Information Theory*, vol. 41, no. 1, pp. 171–187, 1995.
- [138] ——, “A simple recursive algorithm for diagnosis of abrupt changes in random signals,” *IEEE Transactions on Information Theory*, vol. 46, no. 7, pp. 2740–2746, 2000.
- [139] ——, “A lower bound for the detection/isolation delay in a class of sequential tests,” *IEEE Transactions on Information Theory*, vol. 49, no. 11, pp. 3037–3047, 2003.
- [140] S. Nitinawarat and V. V. Veeravalli, “Universal quickest outlier detection and isolation,” in *Proceedings of the IEEE International Symposium on Information Theory (ISIT)*, 2015, pp. 770–774.
- [141] ——, “Universal scheme for optimal search and stop,” *Bernoulli*, vol. 23, no. 3, pp. 1759–1783, 2017.
- [142] Nurjahan, F. Nizam, S. Chaki, S. Al Mamun, and M. S. Kaiser, “Attack detection and prevention in the cyber physical system,” in *Proceedings of the International Conference on Computer Communication and Informatics (ICCCI)*, 2016, pp. 1–6.
- [143] R. J. Oakland and J. S. Oakland, *Statistical Process Control*. Routledge, 2018.
- [144] Y. Ogata, “Statistical models for earthquake occurrences and residual analysis for point processes,” *Journal of the American Statistical Association*, vol. 83, no. 401, pp. 9–27, 1988.
- [145] T. Omi, Y. Ogata, Y. Hirata, and K. Aihara, “Forecasting large aftershocks within one day after the main shock,” *Scientific Reports*, vol. 3, 2013.
- [146] T. Oskiper and H. V. Poor, “Online activity detection in a multiuser environment using the matrix CUSUM algorithm,” *IEEE Transactions on Information Theory*, vol. 48, no. 2, pp. 477–493, 2002.
- [147] S. Padakandla and S. Bhatnagar, “Reinforcement learning in non-stationary environments,” *arXiv:1905.03970*, 2019.
- [148] O. H. M. Padilla, Y. Yu, D. Wang, and A. Rinaldo, “Optimal non-parametric change point detection and localization,” *arXiv:1905.10019*, 2019.
- [149] E. S. Page, “Continuous inspection schemes,” *Biometrika*, vol. 41, no. 1/2, pp. 100–115, 1954.

- [150] M. Pawlak and A. Steland, "Nonparametric sequential signal change detection under dependent noise," *IEEE Transactions on Information Theory*, vol. 59, no. 6, pp. 3514–3531, 2013.
- [151] M. G. Pedersen and M. Meneghini, "Data-driven estimation of change points reveals correlation between face mask use and accelerated curtailment of the COVID-19 epidemic in Italy," *medRxiv*, 2020.
- [152] L. Peel and A. Clauset, "Detecting change points in the large-scale structure of evolving networks," in *Proceedings of the AAAI Conference on Artificial Intelligence*, 2015, pp. 2914–2920.
- [153] J. C. L. Pinto, T. Chahed, and E. Altman, "Trend detection in social networks using Hawkes processes," in *Proceedings of the IEEE/ACM International Conference on Advances in Social Networks Analysis and Mining*, 2015, pp. 1441–1448.
- [154] M. Pollak, "Optimal detection of a change in distribution," *Annals of Statistics*, vol. 13, no. 1, pp. 206–227, 1985.
- [155] M. Pollak and A. M. Krieger, "Shewhart revisited," *Sequential Analysis*, vol. 32, no. 2, pp. 230–242, 2013.
- [156] M. Pollak and A. G. Tartakovsky, "Asymptotic exponentiality of the distribution of first exit times for a class of Markov processes with applications to quickest change detection," *Theory of Probability & Its Applications*, vol. 53, no. 3, pp. 430–442, 2009.
- [157] A. S. Polunchenko, G. Sokolov, and A. G. Tartakovsky, "Optimal design and analysis of the exponentially weighted moving average chart for exponential data," *Sri Lankan Journal of Applied Statistics, Special Issue: Modern Statistical Methodologies in the Cutting Edge of Science*, vol. 15, no. 4, pp. 57–80, 2014.
- [158] A. S. Polunchenko and A. G. Tartakovsky, "On optimality of the Shiryaev–Roberts procedure for detecting a change in distribution," *Annals of Statistics*, vol. 38, no. 6, pp. 3445–3457, 2010.
- [159] —, "State-of-the-art in sequential change-point detection," *Methodology and Computing in Applied Probability*, vol. 14, no. 3, pp. 649–684, 2012.
- [160] H. V. Poor and O. Hadjiliadis, *Quickest Detection*. Cambridge University Press, 2008.
- [161] H. V. Poor, "Quickest detection with exponential penalty for delay," *Annals of Statistics*, vol. 26, no. 6, pp. 2179–2205, 1998.
- [162] M. Qu, F. Y. Shih, J. Jing, and H. Wang, "Automatic solar filament detection using image processing techniques," *Solar Physics*, vol. 228, no. 1-2, pp. 119–135, 2005.
- [163] R. J. Radke, S. Andra, O. Al-Kofahi, and B. Roysam, "Image change detection algorithms: A systematic survey," *IEEE Transactions on Image Processing*, vol. 14, no. 3, pp. 294–307, 2005.
- [164] V. Raghavan and V. V. Veeravalli, "Quickest change detection of a Markov process across a sensor array," *IEEE Transactions on Information Theory*, vol. 56, no. 4, pp. 1961–1981, 2010.
- [165] M. Raginsky, R. Willett, C. Horn, J. Silva, and R. Marcia, "Sequential anomaly detection in the presence of noise and limited feedback," *IEEE Transactions on Information Theory*, vol. 58, no. 8, pp. 5544–5562, 2012.
- [166] R. Rajagopal, X. Nguyen, S. C. Ergen, and P. Varaiya, "Distributed online simultaneous fault detection for multiple sensors," in *Proceedings of the International Conference on Information Processing in Sensor Networks (IPSN)*, 2008, pp. 133–144.
- [167] R. Rajkumar, I. Lee, L. Sha, and J. Stankovic, "Cyber-physical systems: The next computing revolution," in *Proceedings of the Design Automation Conference*, 2010, pp. 731–736.
- [168] A. Reinhart, "A review of self-exciting spatio-temporal point processes and their applications," *Statistical Science*, vol. 33, no. 3, pp. 299–318, 2018.
- [169] G. Rigaill, "Pruned dynamic programming for optimal multiple change-point detection," *arXiv:1004.0887*, 2010.
- [170] Y. Ritov, "Decision theoretic optimality of the CUSUM procedure," *Annals of Statistics*, vol. 18, no. 3, pp. 1464–1469, 1990.
- [171] G. Rovatsos, J. Jiang, A. D. Domínguez-García, and V. V. Veeravalli, "Statistical power system line outage detection under transient dynamics," *IEEE Transactions on Signal Processing*, vol. 65, no. 11, pp. 2787–2797, 2017.
- [172] G. Rovatsos, V. V. Veeravalli, and G. V. Moustakides, "Quickest detection of a dynamic anomaly in a heterogeneous sensor network," in *2020 IEEE International Symposium on Information Theory (ISIT)*, 2020, pp. 1171–1176.
- [173] G. Rovatsos, S. Zou, A. D. Domínguez-García, and V. V. Veeravalli, "Stochastic control of power supply in data centers," in *Proceedings of the Asilomar Conference on Signals, Systems, and Computers*, 2018, pp. 1470–1474.
- [174] G. Rovatsos, S. Zou, and V. V. Veeravalli, "Quickest change detection under transient dynamics," in *Proceedings of the IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2017, pp. 4785–4789.
- [175] G. Rovatsos, G. V. Moustakides, and V. V. Veeravalli, "Quickest detection of moving anomalies in sensor networks," *arXiv:2007.14475*, 2020.
- [176] G. Rovatsos, S. Zou, and V. V. Veeravalli, "Sequential algorithms for moving anomaly detection in networks," *Sequential Analysis*, vol. 39, no. 1, pp. 6–31, 2020.
- [177] S. Saurav, P. Malhotra, V. TV, N. Gugulothu, L. Vig, P. Agarwal, and G. Shroff, "Online anomaly detection with concept drift adaptation using recurrent neural networks," in *Proceedings of the ACM India Joint International Conference on Data Science and Management of Data*, 2018, p. 78–87.
- [178] J. Sharpnack, A. Singh, and A. Rinaldo, "Changepoint detection over graphs with the spectral scan statistic," in *Proceedings of the International Conference on Artificial Intelligence and Statistics (AISTATS)*, vol. 31, 2013, pp. 545–553.
- [179] J. Shen and N. Zhang, "Change-point model on nonhomogeneous Poisson processes with application in copy number profiling by next-generation DNA sequencing," *Annals of Applied Statistics*, vol. 6, no. 2, pp. 476–496, 2012.
- [180] W. A. Shewhart, "The application of statistics as an aid in maintaining quality of a manufactured product," *Journal of the American Statistical Association*, vol. 20, no. 152, pp. 546–548, 1925.
- [181] —, *Economic Control of Quality of Manufactured Product*. American Society for Quality Control, 1931.
- [182] J. Shi, *Stream of Variation Modeling and Analysis for Multistage Manufacturing Processes*. CRC press, 2006.
- [183] A. N. Shiryaev, "On optimum methods in quickest detection problems," *Theory of Probability & Its Applications*, vol. 8, no. 1, pp. 22–46, 1963.
- [184] —, *Optimal Stopping Rules*. New York: Springer-Verlag, 1978.
- [185] D. O. Siegmund, *Sequential Analysis: Tests and Confidence Intervals*, ser. Springer Series in Statistics. Springer, 1985.
- [186] D. Siegmund, "Change-points: From sequential detection to biology and back," *Sequential Analysis*, vol. 32, no. 1, pp. 2–14, 2013.
- [187] D. Siegmund and B. Yakir, "Detecting the emergence of a signal in a noisy image," *Statistics and Its Interface*, vol. 1, no. 1, pp. 3–12, 2008.
- [188] D. Siegmund, B. Yakir, and N. R. Zhang, "Detecting simultaneous variant intervals in aligned sequences," *Annals of Applied Statistics*, vol. 5, no. 2A, pp. 645–668, 2011.
- [189] D. Siegmund and B. Yakir, "Minimax optimality of the Shiryaev–Roberts change-point detection rule," *Journal of Statistical Planning and Inference*, vol. 138, no. 9, pp. 2815–2825, 2008.
- [190] A. Singh, R. Nowak, and P. Ramanathan, "Active learning for adaptive mobile sensing networks," in *Proceedings of the International Conference on Information Processing in Sensor Networks*, 2006, pp. 60–68.
- [191] Y. S. Soh and V. Chandrasekaran, "High-dimensional change-point estimation: Combining filtering with convex optimization," *Applied and Computational Harmonic Analysis*, vol. 43, no. 1, pp. 122–147, 2017.
- [192] M. Sugiyama, T. Suzuki, and T. Kanamori, *Density Ratio Estimation in Machine Learning*. Cambridge University Press, 2012.
- [193] A. Tajer and H. V. Poor, "Quick search for rare events," *IEEE Transactions on Information Theory*, vol. 59, no. 7, pp. 4462–4481, 2013.
- [194] A. Tartakovsky, I. Nikiforov, and M. Basseville, *Sequential Analysis: Hypothesis Testing and Changepoint Detection*. ser. Monographs on Statistics and Applied Probability 136. Boca Raton, London, New York: Chapman & Hall/CRC Press, Taylor & Francis Group, 2015.
- [195] A. G. Tartakovsky, "On asymptotic optimality in sequential changepoint detection: Non-iid case," *IEEE Transactions on Information Theory*, vol. 63, no. 6, pp. 3433–3450, 2017.
- [196] —, *Sequential Change Detection and Hypothesis Testing: General Non-iid Stochastic Models and Asymptotically Optimal Rules*. ser. Monographs on Statistics and Applied Probability 165. Boca Raton, London, New York: Chapman & Hall/CRC Press, Taylor & Francis Group, 2020.
- [197] A. G. Tartakovsky, M. Pollak, and A. Polunchenko, "Third-order asymptotic optimality of the generalized Shiryaev–Roberts changepoint detection procedures," *Theory of Probability & Its Applications*, vol. 56, no. 3, pp. 457–484, 2012.
- [198] A. G. Tartakovsky, "Multidecision quickest change-point detection: Previous achievements and open problems," *Sequential Analysis*, vol. 27, no. 2, pp. 201–231, 2008.

- [199] ——, “Rapid detection of attacks in computer networks by quickest changepoint detection methods,” in *Data Analysis for Network Cyber-Security*. World Scientific, 2014, pp. 33–70.
- [200] ——, “Asymptotically optimal quickest change detection in multi-stream data—part 1: General stochastic models,” *Methodology and Computing in Applied Probability*, vol. 21, no. 4, pp. 1303–1336, 2019.
- [201] A. G. Tartakovsky and G. V. Moustakides, “State-of-the-art in Bayesian changepoint detection,” *Sequential Analysis*, vol. 29, no. 2, pp. 125–145, 2010.
- [202] A. G. Tartakovsky, A. S. Polunchenko, and G. Sokolov, “Efficient computer network anomaly detection by changepoint detection methods,” *IEEE Journal of Selected Topics in Signal Processing*, vol. 7, no. 1, pp. 4–11, 2012.
- [203] A. G. Tartakovsky, B. L. Rozovskii, R. B. Blažek, and H. Kim, “Detection of intrusions in information systems by sequential changepoint methods,” *Statistical Methodology*, vol. 3, no. 3, pp. 252–293, 2006.
- [204] A. G. Tartakovsky, B. L. Rozovskii, R. B. Blažek, and H. Kim, “A novel approach to detection of intrusions in computer networks via adaptive sequential and batch-sequential change-point detection methods,” *IEEE Transactions on Signal Processing*, vol. 54, no. 9, pp. 3372–3382, 2006.
- [205] A. G. Tartakovsky and V. V. Veeravalli, “Change-point detection in multichannel and distributed systems,” *Applied Sequential Methodologies: Real-World Examples with Data Analysis*, vol. 173, pp. 339–370, 2004.
- [206] ——, “General asymptotic Bayesian theory of quickest change detection,” *Theory of Probability & Its Applications*, vol. 49, no. 3, pp. 458–497, 2005.
- [207] ——, “Asymptotically optimal quickest change detection in distributed sensor systems,” *Sequential Analysis*, vol. 27, no. 4, pp. 441–475, 2008.
- [208] M. K. Titsias, J. Sygnowski, and Y. Chen, “Sequential changepoint detection in neural networks with checkpoints,” *arXiv:2010.03053*, 2020.
- [209] I. M. Toke, “Market making behavior in an order book model and its impact on bid-ask spread,” *arXiv:1003.3796*, 2010.
- [210] C. Truong, L. Oudre, and N. Vayatis, “Selective review of offline change point detection methods,” *Signal Processing*, vol. 167, 2020.
- [211] G. Tsechpenakis, D. N. Metaxas, C. Neidle, and O. Hadjiliadis, “Robust online change-point detection in video sequences,” in *Proceedings of the Conference on Computer Vision and Pattern Recognition Workshops*, 2006, pp. 155–161.
- [212] J. N. Tsitsiklis, “Decentralized detection,” *Advances in Statistical Signal Processing*, vol. 2, pp. 297–344, 1993.
- [213] J. Unnikrishnan, V. V. Veeravalli, and S. P. Meyn, “Minimax robust quickest change detection,” *IEEE Transactions on Information Theory*, vol. 57, no. 3, pp. 1604–1614, 2011.
- [214] P. K. Varshney, *Distributed detection and data fusion*. New York: Springer Verlag, 1996.
- [215] V. V. Veeravalli and T. Banerjee, “Quickest change detection,” *Academic Press Library in Signal Processing: Array and Statistical Signal Processing*, vol. 3, pp. 209–256, 2013.
- [216] V. V. Veeravalli and P. K. Varshney, “Distributed inference in wireless sensor networks,” *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences*, vol. 370, no. 1958, pp. 100–117, Jan. 2012.
- [217] V. V. Veeravalli, “Decentralized quickest change detection,” *IEEE Transactions on Information Theory*, vol. 47, no. 4, pp. 1657–1665, 2001.
- [218] V. V. Veeravalli, T. Basar, and H. V. Poor, “Minimax robust decentralized detection,” *IEEE Transactions on Information Theory*, vol. 40, no. 1, pp. 35–40, 1994.
- [219] D. Vere-Jones, “Stochastic models for earth-quake occurrence,” *Journal of Royal Statistical Society*, vol. 32, no. 1, pp. 1–62, 1970.
- [220] L. Y. Vostrikova, “Detecting “disorder” in multidimensional random processes,” in *Doklady Akademii Nauk*, vol. 259, no. 2, 1981, pp. 270–274.
- [221] D. Wang, K. Lin, and R. Willett, “Statistically and computationally efficient change point localization in regression settings,” *arXiv:1906.11364*, 2019.
- [222] D. Wang, Y. Yu, and A. Rinaldo, “Optimal covariance change point detection in high dimension,” *arXiv:1712.09912*, 2017.
- [223] ——, “Optimal change point detection and localization in sparse dynamic networks,” *arXiv:1809.09602*, 2018.
- [224] ——, “Univariate mean change point detection: Penalization, CUSUM and optimality,” *Electronic Journal of Statistics*, vol. 14, no. 1, pp. 1917–1961, 2020.
- [225] D. Wang, Y. Yu, A. Rinaldo, and R. Willett, “Localizing changes in high-dimensional vector autoregressive processes,” *arXiv:1909.06359*, 2019.
- [226] D. Wang, Y. Yu, and R. Willett, “Detecting abrupt changes in high-dimensional self-exciting Poisson processes,” *arXiv:2006.03572*, 2020.
- [227] Y. Wang and Y. Mei, “Large-scale multi-stream quickest change detection via shrinkage post-change estimation,” *IEEE Transactions on Information Theory*, vol. 61, no. 12, pp. 6926–6938, 2015.
- [228] T. Wieland, “A phenomenological approach to assessing the effectiveness of COVID-19 related nonpharmaceutical interventions in Germany,” *Safety Science*, vol. 131, 2020.
- [229] A. Willsky and H. Jones, “A generalized likelihood ratio approach to the detection and estimation of jumps in linear systems,” *IEEE Transactions on Automatic Control*, vol. 21, no. 1, pp. 108–112, 1976.
- [230] M. Woodroffe, *Nonlinear Renewal Theory in Sequential Analysis*, ser. CBMS-NSF regional conference series in applied mathematics. SIAM, 1982.
- [231] L. Xie, Y. Xie, and G. V. Moustakides, “Asynchronous multi-sensor change-point detection for seismic tremors,” in *Proceedings of the IEEE International Symposium on Information Theory (ISIT)*, 2019, pp. 787–791.
- [232] ——, “Sequential subspace change point detection,” *Sequential Analysis*, vol. 39, no. 3, pp. 307–335, 2020.
- [233] Y. Xie, J. Huang, and R. Willett, “Multiscale online tracking of manifolds,” in *IEEE Statistical Signal Processing Workshop (SSP)*, 2012, pp. 620–623.
- [234] ——, “Change-point detection for high-dimensional time series with missing data,” *IEEE Journal of Selected Topics in Signal Processing*, vol. 7, no. 1, pp. 12–27, 2013.
- [235] Y. Xie and D. Siegmund, “Spectrum opportunity detection with weak and correlated signals,” in *In Proceedings of the Asilomar Conference on Signals, Systems and Computers*, 2012, pp. 128–132.
- [236] ——, “Sequential multi-sensor change-point detection,” *Annals of Statistics*, vol. 41, no. 2, pp. 670–692, 2013.
- [237] Y. Xie, M. Wang, and A. Thompson, “Sketching for sequential change-point detection,” in *Proceedings of the IEEE Global Conference on Signal and Information Processing (GlobalSIP)*, 2015, pp. 78–82.
- [238] B. Yakir, *Extremes in Random Fields: A Theory and Its Applications*. John Wiley & Sons, 2013.
- [239] Y.-C. Yao, “Estimating the number of change-points via Schwarz’s criterion,” *Statistics & Probability Letters*, vol. 6, no. 3, pp. 181–189, 1988.
- [240] Y.-C. Yao and S. T. Au, “Least-squares estimation of a step function,” *Sankhyā: The Indian Journal of Statistics, Series A (1961–2002)*, vol. 51, no. 3, pp. 370–381, 1989.
- [241] Y. Yilmaz, G. V. Moustakides, and X. Wang, “Channel-aware decentralized detection via level-triggered sampling,” *IEEE Transactions on Signal Processing*, vol. 61, no. 2, pp. 300–315, 2013.
- [242] C. E. Yoon, O. O’Reilly, K. J. Bergen, and G. C. Beroza, “Earthquake detection through computationally efficient similarity search,” *Science Advances*, vol. 1, no. 11, p. e1501057, 2015.
- [243] J. Y. Yu and S. Mannor, “Piecewise-stationary bandit problems with side observations,” in *Proceedings of the International Conference on Machine Learning (ICML)*, 2009, p. 1177–1184.
- [244] X. Yu, M. Baron, and P. K. Choudhary, “Change-point detection in binomial thinning processes, with applications in epidemiology,” *Sequential Analysis*, vol. 32, no. 3, pp. 350–367, 2013.
- [245] D. Zhang and L. Balzano, “Global convergence of a Grassmannian gradient descent algorithm for subspace estimation,” in *Proceedings of the International Conference on Artificial Intelligence and Statistics (AISTATS)*, 2016, pp. 1460–1468.
- [246] N. R. Zhang, B. Yakir, C. L. Xia, and D. O. Siegmund, “Scanning a Poisson random field for local signals,” *Annals of Applied Statistics*, vol. 10, no. 2, p. 726, 2016.
- [247] R. Zhang and Y. Mei, “Asymptotic statistical properties of communication-efficient quickest detection schemes in sensor networks,” *Sequential Analysis*, vol. 37, no. 3, pp. 375–396, 2018.
- [248] R. Zhang, Y. Mei, and J. Shi, “Robust real-time monitoring of high-dimensional data streams,” *arXiv:1906.02265*, 2019.
- [249] S. Zhu, A. Bukharin, L. Xie, M. Santillana, S. Yang, and Y. Xie, “High-resolution spatio-temporal model for county-level COVID-19 activity in the US,” *arXiv:2009.07356*, 2020.
- [250] S. Zhu, R. Ding, M. Zhang, P. Van Hentenryck, and Y. Xie, “Spatio-temporal point processes with attention for traffic congestion event modeling,” *arXiv:2005.08665*, 2020.

- [251] S. Zou, G. Fellouris, and V. V. Veeravalli, "Asymptotic optimality of D-CuSum for quickest change detection under transient dynamics," in *Proceedings of the IEEE International Symposium on Information Theory (ISIT)*, 2017, pp. 156–160.
- [252] S. Zou, Y. Liang, H. V. Poor, and X. Shi, "Nonparametric detection of anomalous data streams," *IEEE Transactions on Signal Processing*, vol. 65, no. 21, pp. 5785–5797, 2017.
- [253] S. Zou, V. V. Veeravalli, J. Li, D. Towsley, and A. Swami, "Distributed quickest detection of significant events in networks," in *Proceedings of the IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2019, pp. 8454–8458.
- [254] S. Zou, Y. Liang, and H. V. Poor, "Nonparametric detection of geometric structures over networks," *IEEE Transactions on Signal Processing*, vol. 65, no. 19, pp. 5034–5046, 2017.
- [255] S. Zou, V. V. Veeravalli, J. Li, and D. Towsley, "Quickest detection of dynamic events in networks," *IEEE Transactions on Information Theory*, vol. 66, no. 4, pp. 2280–2295, 2019.



Yao Xie (S'04-M'12) received the B.Sc. degree in electrical engineering and computer science from the University of Science and Technology of China, the M.Sc. degree in electrical and computer engineering from the University of Florida, and the Ph.D. degree in electrical engineering (minor in Mathematics) from Stanford University. She is an Associate Professor and Harold R. and Mary Anne Nash Early Career Professor at Georgia Institute of Technology in the H. Milton Stewart School of Industrial and Systems Engineering, and an Associate Director of

the Machine Learning Center. She was a Research Scientist at Duke University. Her research areas are statistics (particularly change-point detection and spatio-temporal data modeling), machine learning, and signal processing, providing the theoretical foundation and developing computationally efficient and statistically powerful algorithms for engineering problems. She received the National Science Foundation (NSF) CAREER Award in 2017. She is currently an Associate Editor for *IEEE Transactions on Signal Processing*, *Sequential Analysis: Design Methods and Applications*, *INFORMS Journal on Data Science*, and serves on the Editorial Board of *Journal of Machine Learning Research*.



Liyan Xie (S'19) received the B.Sc. degree in statistics from the University of Science and Technology, China, in 2016. She is currently pursuing the Ph.D. degree with the H. Milton Stewart School of Industrial and Systems Engineering, Georgia Institute of Technology. She is also 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. Her research interests are the intersection between statistics and optimization, with a primary focus on sequential change detection and their applications in sensor networks and healthcare. She was awarded the finalist of the INFORMS QSR Student Paper Competition and Runner-up for the INFORMS Computing Society Student Paper Prize in 2019.

primary focus on sequential change detection and their applications in sensor networks and healthcare. She was awarded the finalist of the INFORMS QSR Student Paper Competition and Runner-up for the INFORMS Computing Society Student Paper Prize in 2019.



Venugopal V. Veeravalli (M'92-SM'98-F'06) received the B.Tech. degree (Silver Medal Honors) from the Indian Institute of Technology, Bombay, in 1985, the M.S. degree from Carnegie Mellon University, Pittsburgh, PA, in 1987, and the Ph.D. degree from the University of Illinois at Urbana-Champaign, in 1992, all in electrical engineering.

He joined the University of Illinois at Urbana-Champaign in 2000, where he is currently the Henry Magnuski Professor in the Department of Electrical and Computer Engineering, and where he is also affiliated with the Department of Statistics, and the Coordinated Science Laboratory. Prior to joining the University of Illinois, he was on faculty of the ECE Department at Cornell University. He served as a Program Director for communications research at the U.S. National Science Foundation from 2003 to 2005. His research interests span the theoretical areas of statistical inference, machine learning, and information theory, with applications to data science, wireless communications, and sensor networks.

Prof. Veeravalli was a Distinguished Lecturer for the IEEE Signal Processing Society during 2010–2011. He has been on the Board of Governors of the IEEE Information Theory Society. He has been an Associate Editor for Detection and Estimation for the *IEEE Transactions on Information Theory* and for the *IEEE Transactions on Wireless Communications*. Among the awards he has received for research and teaching are the IEEE Browder J. Thompson Best Paper Award, the Presidential Early Career Award for Scientists and Engineers (PECASE), and the Wald Prize in Sequential Analysis.



Shaofeng Zou (S'14-M'16) received the Ph.D. degree in Electrical and Computer Engineering from Syracuse University in 2016. He received the B.E. degree (with honors) from Shanghai Jiao Tong University, Shanghai, China, in 2011. He was a postdoctoral research associate at the Coordinated Science Lab, University of Illinois at Urbana-Champaign during 2016–2018. He joined the Department of Electrical Engineering, University at Buffalo, the State University of New York in 2018, where he is currently an Assistant Professor. Dr. Zou's research interests include statistical signal processing, machine learning, and information theory.

interests include statistical signal processing, machine learning, and information theory.