

Robust Change Detection via Information Projection

Deniz Sargun, C. Emre Koksal, *Senior Member, IEEE*

Abstract

We study the robust transient and quickest change detection problems with unknown post-change distributions. When the distribution after the change point is unknown, the change in the distribution of observations may occur in multiple ways without much structure on the observations, whereas, before the change point, a false alarm is highly structured, following a particular sample path. We first characterize these likely events for the deviation and propose a method to test the empirical distribution, relative to the most likely way for it to occur as an outlier. We benchmark our method with the finite moving average (FMA) method, generalized likelihood ratio test (GLRT) and M-statistic kernel (MSK) change point detection method under 4 different performance criteria including the run time complexity. Finally, we apply our method on economic market indicators and climate data. Our method successfully captures the regime shifts during times of historical significance for the markets and identifies the current climate change phenomenon to be a highly likely regime shift rather than a random event.

Index Terms

composite hypothesis testing, quickest change detection, transient change detection, unknown post-change distribution, time complexity, KL divergence, information projection

I. INTRODUCTION

Not every long-term deviation from the norm should be considered as a result of a change. Such occurrences may also be caused by rare events driven by the system and this distinction is important in many applications from model selection [1], [2] to hardware faults [3], [4], security [5], [6] and even health [7], [8]. To that end, this paper focuses on problem instances of change detection [9] (also see [10] for a wide treatment of change detection and refer to chapter 8 for nonadditive change models and existing approaches) in a time-series with minimax cost with respect to change point and post-change distributions.

The authors are with the Department of Electrical and Computer Engineering, The Ohio State University, Columbus, OH 43210 USA (e-mail:{sargun.1, koksal.2}@osu.edu).

The quickest change detection problem with unknown post-change models has been studied widely, both in theory [11]–[16] and in application [17], [18] (see [19] for more variants and applications). For some non-Bayesian quickest change detection problems where pre- and post-change distributions have finitely many alternatives it is known that a version of the cumulative sum algorithm (CUSUM) is optimal or asymptotically optimal [20]–[22]. For other problems where the pre- or post-change distribution has infinitely many alternatives or is not parametrized, algorithms that form a likelihood ratio (ex. CUSUM) are not applicable. In this case, since the change in the distribution is arbitrary it is hard to perform optimally against all alternatives. A widely used method in the literature is to form a generalized likelihood ratio statistics over the set of alternative distributions and then use it as a drift term [14], [21], [23], which can be computationally demanding or without any optimality guarantees. To address this, we propose a recursive¹ and linear complexity online method, called the information projection test (IPT), that uses an initial convex decision region and modifies it to suppress most likely false alarms. IPT utilizes the separation between pre- and post-change distributions at the first level and compares the relative entropy of the empirical distribution with respect to the distribution of the most likely outlier at the second level. For the quickest change detection problem, we find bounds on the average run length (ARL) and the worst average detection delay (WADD), prove its asymptotic optimality up to a multiplicative constant and share simulation results which empirically show that IPT performs similar to GLRT.

To motivate our approach, we first define IPT for fixed window size tests in Section IV and bound its receiver operating characteristic (ROC) for the composite hypothesis testing problem in Subsection V-A. Then, we transition to the robust transient change detection problem to introduce the concept of delay and bound IPT’s performance of detecting changes with bounded delay in Subsection V-B. Next, we prove IPT is asymptotically optimal up to a constant for a robust version of Lorden’s criteria in Subsection V-C. We provide empirical comparisons of detection performance with FMA [24], GLRT and MSK [16] in Section VI.

Another important metric considered is the computational complexity of the algorithms in Section VII. There is already research on low complexity change point detection some examples of which are detecting changes in video content popularity [25] and detecting propagating changes in networks [26], [27]. Even with a similar detection performance, IPT has a few orders of lower complexity (99.86% reduction in test time for an alphabet with ~ 6500 letters and same order of samples) compared to GLRT. Thus, we claim

¹if the function q can be estimated recursively for empirical probability mass function \hat{f}

our linear complexity algorithm achieves a better trade-off between computation time and performance compared to other methods in the literature.

Contributions of this paper can be summarized as follows:

- We introduce a simple and novel algorithm for detecting changes with known pre-change distribution but unknown post-change distribution.
- We propose a novel robust transient change detection problem.
- We prove performance bounds in the non-asymptotic regime and optimality up to a multiplicative constant in the asymptotic regime.
- We empirically benchmark our method for 3 problems: composite hypothesis testing problem, robust transient change detection problem and quickest change detection problem with robust Lorden's criteria.
- We propose widely applicable scenarios where our problem assumptions hold and study special cases to provide insight for our method.
- We compare the time complexity of our approach against existing methods theoretically and empirically.
- We draw new insights by applying our idea in analyzing economic and climate time series data to understand whether periods of long deviations from the norm are outliers or change. We identify shifts from the average behavior of company returns and market indices to show that our scheme successfully manages to detect the periods with regime changes like crises. One of the highlights of our analysis is that the climate change phenomenon that is observed in the last 30 years is highly unlikely to be an outlier, giving credence to the hypothesis that it is caused by exogenous factors.

Earlier version of this work has appeared in [28]. In this version we include performance bounds for the transient and quickest change detection problems, include empirical results for the latter, increase the number of methods we compare against and applications on real datasets, provide application scenarios and study insightful special cases.

II. MODEL AND PROBLEM STATEMENT

Let $\mathcal{A} = \{a_1, \dots, a_m\}$ be a finite alphabet and \mathcal{P} denote the probability simplex of probability mass functions (p.m.f.s) f over \mathcal{A} equipped with the l_1 norm. Consider $X_1^\infty = (X_1, \dots, X_k, \dots)$, a sequence of observations from \mathcal{A} , where, given change point t_1 , each random variable in $X_1^{t_1-1} = (X_1, \dots, X_{t_1-1})$

is independent and identically distributed (i.i.d.) with known pre-change p.m.f. f_0 and $X_{t_1}^\infty$ are i.i.d. with the unknown post-change p.m.f. $f_1 \in \mathcal{P}_1$, independent of previous observations $X_1^{t_1-1}$. Assume there exists $q: \mathcal{P} \rightarrow \mathbb{R}$ quasiconcave and Lipschitz continuous with Lipschitz constant L , that satisfies $q(f_0) = q_0 < 0 < q \leq q_1 = q(f_1)$. In other words, \mathcal{P}_1 is a subset of a closed and convex set $q^{-1}([q, \infty))$ and $f_0 \notin q^{-1}([q, \infty))$. We use $\hat{f}_{X_i^j}$ for the empirical p.m.f. of X_i^j , i.e. $\hat{f}_{X_i^j}(x) = \frac{1}{j-i+1} \sum_{k=i}^j \mathbb{1}_{x_i}(x)$, and \mathcal{P}^n for the discrete set of empirical p.m.f.s that can be realized with n samples from \mathcal{A} . Finally, $I(f \| f') = \sum_a f(a) \log \frac{f(a)}{f'(a)}$ denotes the Kullback-Leibler (KL) divergence between distributions f and f' in nats and P_{f_1, t_1} and E_{f_1, t_1} denote the probability law and expectation when post-change distribution is f_1 and change time is t_1 respectively. When change does not occur we denote $P_{f_1, \infty}$ and E_{f_1, t_1} briefly as P_∞ and E_∞ respectively.

First, we consider the composite hypothesis testing problem. Given n , the null hypothesis, \mathbf{H}_0 , is true if X_1^n are i.i.d. f_0 and the alternative hypothesis, \mathbf{H}_1 , is true if i.i.d. f_1 . Our goal is to minimize the maximum probability of misdetection given an upper bound on the probability of false alarm. Since samples are i.i.d. under both hypotheses, the empirical p.m.f. $\hat{f}_{X_1^n}$ provides a sufficient statistic and if we define the decision regions $\Gamma_1 = \{\hat{\mathbf{H}} = \mathbf{H}_1\}$ and $\Gamma_0 = \mathcal{P}^n \setminus \Gamma_1 = \{\hat{\mathbf{H}} = \mathbf{H}_0\}$ the problem can be stated as

$$\inf_{\Gamma} \sup_{f_1} P_{f_1, 1}(\hat{f}_{X_1^n} \in \Gamma_0) \quad (1)$$

$$\text{subject to } P_\infty(\hat{f}_{X_1^n} \in \Gamma_1) \leq \alpha \quad (2)$$

where $\Gamma = (\Gamma_0, \Gamma_1)$. The decision rule

$$\hat{\mathbf{H}}(f) = \begin{cases} \mathbf{H}_1, & q(f) \geq c^S \\ \mathbf{H}_0, & \text{otherwise} \end{cases}$$

is given as an example in Fig. 1a for alphabet size $m = 3$.

Second, we consider the transient change detection problem [22] with unknown post-change distribution. The robust transient change detection problem minimizes the probability of misdetecting a single transient change within a specified window after the change. Without loss of generality (w.l.o.g.), we assume the change in distribution is permanent but the change point t_1 has to be detected in a bounded time n which is not greater than the duration of the change in the initial problem². The minimization is subject to an upper bound on the probability of false alarm within any window of size n_α . If t_a denotes an arbitrary

²When the detection window, n in (3), is assumed to be smaller than the duration of transient change, the distribution of samples after the transient change becomes an invariant of the probabilities of detection considered in (3) and (4). Therefore, we may assume the change is permanent.

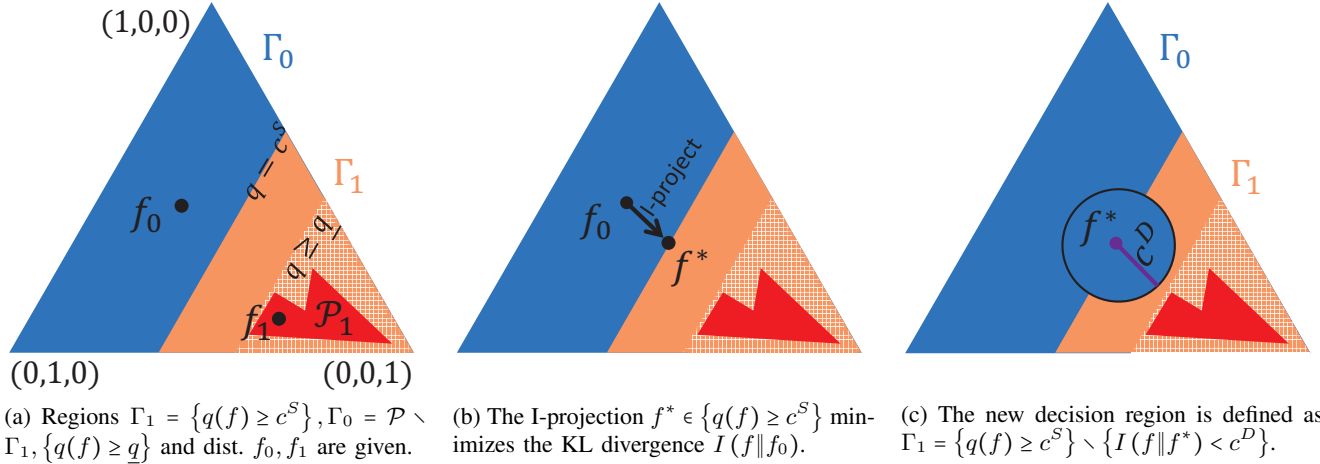


Fig. 1. The probability simplex with $m = 3$ for our setup is illustrated. Each corner corresponds to a different deterministic distribution. The pre-change distribution f_0 and its I-projection f^* onto $\{q(f) \geq c^S\}$ are shown. Decision regions (old and new) Γ_0 and Γ_1 are shown in blue and light orange respectively. The (nonconvex) set of post-change distributions \mathcal{P}_1 is shown in red.

alarm (stopping) time of X_1^∞ , the problem can be formulated as

$$\inf_{t_a} \sup_{f_1, t_1} P_{f_1, t_1} (t_a \geq t_1 + n | t_a \geq t_1) \quad (3)$$

$$\text{subject to } \sup_k P_\infty (k \leq t_a < k + n_\alpha) \leq \tilde{\alpha} \quad (4)$$

where we assume the worst post-change distribution f_1 is realized. The optimal solution to this problem is unknown even for known post-change distribution except for the case $n = 1$ subject to a minimum ARL [22].

Finally, we consider the quickest change detection problem with unknown post-change distribution. The quickest change detection in Lorden's sense [29] involves minimizing WADD given a minimum ARL. In this robust version we assume the worst post-change distribution. Then, the problem can be formulated as

$$\inf_{t_a} \sup_{f_1, t_1} \text{esssup} E_{f_1, t_1} ((t_a - t_1 + 1)^+ | X_1^{t_1-1}) \quad (5)$$

$$\text{subject to } E_\infty t_a \geq \gamma. \quad (6)$$

Optimal solution to this problem is also unknown. Therefore, we compare our solution of (3), (4) and (5), (6) with frequently used (asymptotically optimal) methods in the literature.

III. MOTIVATION

The first problem we consider is an instance of composite hypothesis testing [30] without a parametric model under \mathbf{H}_1 . In this case, a typical detector picks \mathbf{H}_1 if the empirical distribution, $\hat{f} = \hat{f}_{X_1^n}$ is closer to a distribution f_1 in \mathcal{P}_1 than it is to f_0 . Even in that case, however, it is not necessarily true that X_1^n is drawn from some post-change distribution. It may be drawn from f_0 , yet the empirical distribution might look as if it is drawn from some f_1 , leading to a false positive/alarm. We call such strings that raise an alarm but are close to f_0 an outlier.

For clarity, we provide insight from large deviations theory where the main results are asymptotic, but our contributions apply to the non-asymptotic regime as well. We study the way outliers occur using Sanov's theorem and we are motivated by it to detect and suppress such outliers, as they the most likely way for a false alarm to occur. See [31] for a detailed treatment.

Theorem (Sanov). *For any continuous and quasiconcave function $q : \mathcal{P} \rightarrow \mathbb{R}$, p.m.f. f_0 and constant $c \neq \sup_f q(f)$,*

$$\lim_{n \rightarrow \infty} -\frac{1}{n} \log P_\infty (q(\hat{f}_{X_1^n}) \geq c) = \inf_{q(f) \geq c} I(f \| f_0) \quad (7)$$

The proof of the most general form can be found in [32]. Note that the interesting case is when $q(f_0) < c < \sup_f q(f)$ and that we have the equality in Sanov's theorem, since $\{q(f) \geq c\}$ is the closure of its interior as shown in Fig. 1b, where we also illustrate Sanov's theorem. The solution of the constrained convex optimization $\operatorname{argmin}_{q(f) \geq c} I(f \| f_0)$ is called the information projection or I-projection of f_0 onto the superlevel set $\{q(f) \geq c\}$ and is denoted by f^* . The theorem gives an asymptotic result in the number of samples and provides the exact characterization of the exponent at which the probability of false alarm decays with n . The theorem also implies that the most likely way for a sustained deviation to occur under \mathbf{H}_0 is when the empirical distribution of the associated string looks as if it is drawn from the I-projection f^* . The probability of this particular way of deviation dominates the probability of all others asymptotically. An approximate version of this result also holds for a finite string of length n if Stirling's approximation is satisfied. In Section IV, we give our algorithm which is based on using the projected distribution as the most likely outlier distribution. Our algorithm not only suppresses the projection f^* but also any empirical distribution in the proximity of f^* measured in KL divergence. Subsequently, we evaluate the detection performance in the non-asymptotic region to justify its applicability.

A. The choice of function q

One of our main assumptions is the existence of a function q to differentiate pre- and post-change conditions. Below we characterize a few such scenarios where we can define a Lipschitz continuous quasiconcave function to be utilized in change detection to distinguish the unknown post-change distribution from the pre-change distribution.

1) *Change in the Expectation of a Function of the Observations:* In various change detection problems [33], [34], the change is usually manifested in the mean of the observations. More generally, assume there exists a function $h : \mathcal{A} \rightarrow \mathbb{R}$ and $\underline{h} > 0$ such that $E_\infty h(X_k) \leq E_{f_{1,1}} h(X_k) + \underline{h}$. In other words, assume we can quantify the change as an increase (or decrease) in the mean of a function h of the observations. Then, the function $E_f h(X)$ over \mathcal{P} can be used to define a function q that satisfies the conditions in Section II ³. Note that, a change in any moment of the observations can be considered in this category.

2) *Change in Variance:* Similarly, an *increase* in the variance of a distribution also yields a concave and Lipschitz function q over \mathcal{P} .

3) *Discrepancy in Post-Change Conditions:* In certain problems the post-change distribution might be within a certain error of a known distribution. For example, assume there exists f_{1*} such that for all $f_1 \in \mathcal{P}_1$, $\|f_1 - f_{1*}\|_1 \leq \delta$. If $\delta \leq (1 - \rho) \frac{I(f_{1*} \| f_0)}{\sup_a \left| \log \frac{f_{1*}(a)}{f_0(a)} \right|}$ for some $\rho > 0$, the function q defined as $q(f) = I(f \| f_0) - I(f \| f_{1*})$ satisfies the conditions in Section II with $q_0 = -I(f_0 \| f_{1*}) < \rho I(f_{1*} \| f_0) = \underline{q}$.

4) *A Nonconvex \mathcal{P}_1 :* Some quickest change detection problems [35], [36] assume (or pick) a finite set of post-change distributions. More generally, let us assume \mathcal{P}_1 is a nonconvex subset of \mathcal{P} . Further, assume the closure of the convex hull of \mathcal{P}_1 does not include the pre-change distribution f_0 , i.e. $f_0 \notin E = \overline{\text{conv}(\mathcal{P}_1)}$. Then, E is also a convex set and we can use the distance to the set E in \mathcal{P} to construct a function q that satisfies the conditions in Section II ⁴.

IV. INFORMATION PROJECTION TEST

The information projection test (IPT) exploits the observation that outliers occur in a particular way with high probability, while there is no such structure for the deviations.

³The function $\sum_a h(a)f(a)$ is linear in f and has Lipschitz constant $L = \sup_a |h(a)|$.

⁴The distance $\|f - E\|_1$ is convex in f over \mathcal{P} .

A. The Algorithm

We describe the fixed window size IPT for hypothesis testing and transient change detection here and the variable size IPT for the quickest change detection in Subsection V-C. Given a window size n , the fixed window IPT involves the following steps.

- 1) Pick $c^S, c^D \geq 0$ and window size n .
- 2) Find the I-projection $f^* = \operatorname{argmin}_{q(f) \geq c^S} I(f \| f_0)$.
- 3) At time $k \geq n$, find $S_k = q(\hat{f}_{X_{k-n+1}^k})$.
- 4) If $S_k \geq c^S$, find $D_k = I(\hat{f}_{X_{k-n+1}^k} \| f^*)$; else, return to Step 3.
- 5) If $D_k \geq c^D$, stop and claim a change has happened; else, return to Step 3.

Note that Step 5 determines if the empirical distribution \hat{f} is close to the most likely distribution for an outlier, f^* . This process is illustrated in Fig. 1c. In the rest of the paper, we call the relative entropy $I(f \| f^*)$ also as the relative log likelihood function (RLLF) with respect to f^* , referring to the comparison of the likelihood of the projection distribution with that of the empirical distribution of the observation. We also use c to denote (c^S, c^D) .

B. Finding the I-projection

In the most general case, the information projection f^* of a distribution f_0 onto a set B is defined as

$$\begin{aligned} & \inf_f I(f \| f_0) \\ & \text{subject to } f \in B \end{aligned}$$

The KL divergence between f and f_0 is strictly convex in the optimization variable f . Therefore, if the set B is a convex subset of distribution, the I-projection becomes a convex problem.

In this paper, we assume a finite alphabet and a convex set B parametrized with q . Thus, we can simplify the problem as

$$\begin{aligned} & \inf_f \sum_a f(a) \log \frac{f(a)}{f_0(a)} \\ & \text{subject to } q(f) \geq c^S \end{aligned}$$

which can be solved, for example, using gradient descent and projection methods. [37], [38] Further, if q is concave (in contrast to quasiconcave), differentiable and $\{q(f) \geq c^S\}$ has nonempty relative interior,

KKT conditions are necessary and sufficient for f^* [39].

V. PERFORMANCE GUARANTEES

A. Composite Hypothesis Testing

The next two theorems bound the probability of false alarm and misdetection as a function of c . We show that the false positive probability decays exponentially with $c^D + \frac{((c^S - q_0)^+)^2}{2L^2}$ and that the worst case probability of misdetection decays with rate $\left(\frac{q - q_0 - (c^S - q_0)^+}{\sqrt{2}L} - \sqrt{c^D}\right)^2$.

Theorem 1. For any c such that $c^S < \sup_f q(f)$,

$$P_\infty(\hat{f}_{X_1^n} \in \Gamma_1) \leq (n+1)^m \exp\left(-n\left(c^D + \frac{((c^S - q_0)^+)^2}{2L^2}\right)\right). \quad (8)$$

Theorem 2. For any c such that $(c^S - q_0)^+ + \sqrt{2L^2 c^D} < \underline{q} - q_0$,

$$\sup_{f_1} P_{f_1,1}(\hat{f}_{X_1^n} \in \Gamma_0) \leq (n+1)^m \exp\left(-n\left(\frac{q - q_0 - (c^S - q_0)^+}{\sqrt{2}L} - \sqrt{c^D}\right)^2\right). \quad (9)$$

B. Transient Change Detection

In this subsection, we prove corollaries to Theorem 1 and 2 that provide a bound on the transient change detection performance of IPT under the assumption of unknown post-change distribution. The transient change detection problem considers the worst case probability of false alarm and the worst case probability of misdetection within a given length of interval as its metrics. Since we are given a window of n samples to detect the change and $n \ll n_\alpha$, we utilize a fixed window IPT. This problem was first proposed in [40] and is also studied recently in [22], [24].

Corollary 3. For any c such that $(c^S - q_0)^+ + \sqrt{2L^2 c^D} < \underline{q} - q_0$ the IPT stopping time t_I with fixed window length $\frac{n+1}{2}$ satisfies

$$\sup_k P_\infty(k \leq t_I < k + n_\alpha) \leq 1 - (1 - \exp(-\nu n))^{\lceil \frac{2n_\alpha}{n+1} \rceil} \quad (10)$$

$$\sup_{f_1, t_1} P_{f_1, t_1}(t_I \geq t_1 + n | t_I \geq t_1) \leq \exp(-\eta n) \quad (11)$$

where

$$\nu = \frac{c^D}{2} + \frac{((c^S - q_0)^+)^2}{4L^2} - \frac{m \log\left(\frac{n+3}{2}\right)}{n+1} \quad (12)$$

$$\eta = \left(\frac{q - q_0 - (c^S - q_0)^+}{2L} - \sqrt{\frac{c^D}{2}} \right)^2 - \frac{m \log\left(\frac{n+3}{2}\right)}{n+1}. \quad (13)$$

C. Quickest Change Detection

In this subsection, we apply IPT to the quickest change detection problem with an unknown post-change distribution. To minimize the average detection delay under the worst change time t_1 and pre-change realizations $X_1^{t_1-1}$, we modify IPT to have an effective window that only contains the samples with a positive drift from the minimum the random walk has visited, similar to the CUSUM algorithm. We also use a restart mechanism if the empirical distribution is similar to a most likely false alarm, similar to the IPT with fixed window size. Since, at any given time k , the effective window size n_k is varying, we define the most likely deviation as a function of the window size, f_n^* .

$$f_n^* = \operatorname{argmin}_{q(f) \geq \frac{c^S}{n}} I(f \| f_0) \quad (14)$$

For any k , τ_k denotes the last time the IPT has been restarted since the empirical distribution resembles a false alarm. So, $\tau_1 = 1$ and for any k ,

$$(S_k, i_k) = \max_{\tau_k \leq i \leq k+1} (k - i + 1) q(\hat{f}_{X_i^k}) \quad (15)$$

We define the index i_k that maximizes S_k to define the effective window size $n_k = k - i_k + 1$. Finally, we compare the proximity of the empirical distribution over the effective window with respect to the most likely false alarm given the window size n_k . If D_k is large, we stop the observation process and raise an alarm, else we conclude a false alarm and reset the algorithm. The threshold for this step is allowed to vary with n , so we select a sequence of $c_n^D \geq 0$ instead of a constant c^D .

$$D_k = I(\hat{f}_{X_{i_k}^k} \| f_{n_k}^*) \quad (16)$$

$$\tau_{k+1} = \begin{cases} k + 1, & S_k \geq c^S, D_k < c_{n_k}^D \\ \tau_k, & \text{otherwise} \end{cases} \quad (17)$$

$$t_I = \inf \{k | (S_k, D_k) \geq (c^S, c_{n_k}^D)\} \quad (18)$$

We also use shorthand notations c_n and $Q(i, j)$ to refer to $\left(\frac{c^S}{n}, c_n^D\right)$ and $(j-i+1)q\left(\hat{f}_{X_i^j}\right)$ respectively. We define the auxiliary stopping time $t_S = \inf\{k|S_k \geq c^S\}$ where only the first statistic crosses the threshold and we define t_S^k and t_I^k to be the stopping times starting at k , counting only the observations after X_{k-1} .

In the following theorem we show that the ARL of our method can be lower bounded in terms of c . We prove the theorem where q is the expectation operator but the general result can be obtained by constructing a separating hyperplane between the convex and closed set \mathcal{P}_1 and the initial distribution $f_0 \notin \mathcal{P}_1$. Then, step 1 of the IPT holds only if the hyperplane is crossed. Thus, a lower bound on $E_\infty t_S$ can be formulated using this halfspace instead of \mathcal{P}_1 .

Theorem 4. *If $c_n^D = c^D \geq \frac{2|q_0|q}{(1+\rho)L^2}$ for $n > (1+\rho)\frac{c^S}{q}$ for some $\rho > 0$, then, the ARL of IPT is bounded as*

$$E_\infty t_I \geq \frac{\exp(v^* c^S)}{2^m \exp\left(-\frac{2|q_0|}{L^2} c^S\right) + \frac{\left((1+\rho)\frac{c^S}{q} + 1\right)^m \exp\left(-\frac{c^D}{q} c^S\right)}{1 - \left((1+\rho)\frac{c^S}{q} + 1\right)^{\frac{mq}{(1+\rho)c^S}} \exp(-c^D)} \rightarrow \exp\left(\left(v^* + \frac{2|q_0|}{L^2}\right) c^S\right) \quad (19)$$

where $\psi(v) = E_\infty \exp(vX)$ and $v^* > 0$ satisfies $\psi(v^*) = 1$.

The following lemma proves a bound on the stopping time for any change time t_1 and observations $X_1^{t_1-1}$.

Lemma 5. *For any t_1, f_1 and $X_1^{t_1-1}$, the IPT stopping time t_I is bounded by*

$$t_I \leq t_S^{(j)} + t_I^{t_S^{(j)}+1} \quad (20)$$

for any $j = 0, \dots$ where $t_S^{(j)}$ is the j th time S_k crosses c^S , i.e. $t_S^{(0)} = 0$ and $t_S^{(j+1)} = t_S^{(j)} + t_S^{t_S^{(j)}+1}$. Further, for any $j = 1, \dots$ and $t_S^{(j-1)} < k \leq t_S^{(j)}$,

$$t_S^{(j)} \leq k + t_S^k - 1. \quad (21)$$

Thus, for $j^* = \min\{j | t_S^{(j)} \geq t_1\}$,

$$t_I \leq t_1 + t_S^{t_1} - 1 + t_I^{t_S^{(j^*)}+1}. \quad (22)$$

In the next lemma, we prove that the average delay is bounded by a function of the average delay when $t_1 = 1$ using the bound we have found in Lemma 5. This lemma shows that the worst average delay IPT experiences is bounded by the average delay up to a multiplicative constant.

Lemma 6. For any f_1, t_1 and $X_1^{t_1-1}$, the average delay is bounded as follows.

$$E_{f_1, t_1} \left((t_I - t_1 + 1)^+ | X_1^{t_1-1} \right) \leq E_{f_1, 1} (t_S + t_I) \quad (23)$$

Before showing a bound on the WADD, we bound the average detection delay when change happens at $t_1 = 1$ asymptotically as $c^S \rightarrow \infty$.

Lemma 7. For any f_1 , if $c_n^D = 0$ for $n \leq (1 + \rho) \frac{c^S}{\underline{q}}$ for some $\rho > 0$, then, as $c^S \rightarrow \infty$

$$E_{f_1, 1} t_I \leq \frac{c^S}{\underline{q}_1}. \quad (24)$$

The next theorem is the asymptotic characterization of the WADD of IPT. In the proof we first derive a non-asymptotic bound and then consider the limit as $c^S \rightarrow \infty$.

Theorem 8. If $c_n^D = 0$ for $n \leq (1 + \rho) \frac{c^S}{\underline{q}}$ for some $\rho > 0$, then, as $c^S \rightarrow \infty$,

$$WADD(t_I) \leq \frac{2c^S}{\underline{q}} \quad (25)$$

Lorden showed in his seminal work [29] that the optimal WADD for any change detection algorithm with ARL at least γ is $\frac{\log \gamma}{I(f_1 \| f_0)}$ asymptotically as $\gamma \rightarrow \infty$ even if the post-change distribution f_1 is known. Using Theorem 4 and 8, we prove that IPT is asymptotically optimal in robust Lorden's criteria up to a multiplicative constant.

Corollary 9. For $\rho > 0$, $c^D \geq \frac{2|q_0|q}{(1+\rho)L^2}$ and $c_n^D = \begin{cases} 0, & n \leq (1 + \rho) \frac{c^S}{\underline{q}}, \\ c^D, & \text{otherwise} \end{cases}$, as $\gamma \rightarrow \infty$,

$$WADD(t_I) = \frac{\log \gamma}{\underline{q} \left(\frac{v^*}{2} + \frac{|q_0|}{L^2} \right)}. \quad (26)$$

Proof: Use Theorem 4 and 8. ■

We end our discussion on the quickest change detection problem by describing two special cases.

1) *Change in Mean*: For the special case where the change is in the mean, IPT reduces to a CUSUM algorithm with restart.

$$q\left(\hat{f}_{X_i^k}\right) = \frac{1}{k-i+1} \sum_{j=i}^k X_j \quad (27)$$

$$S_k = \max_{\tau_k \leq i \leq k} \sum_{j=i}^k X_j \quad (28)$$

$$= \begin{cases} (S_{k-1} + X_k)^+, & \tau_k \neq k \\ (X_k)^+, & \tau_k = k \end{cases} \quad (29)$$

Further, the most likely deviations f_n^* become the exponentially tilted distributions with mean $\max\left\{\frac{c^S}{n}, q_0\right\}$.

$$f_n^* = \underset{\sum_a a f(a) \geq \frac{c^S}{n}}{\operatorname{argmin}} I(f \| f_0) \quad (30)$$

$$f_n^*(a) = f_0(a) \exp(ra - \Lambda(r)) \quad (31)$$

$$D_k = I(\hat{f} \| f_{n_k}^*) \quad (32)$$

$$= I(\hat{f} \| f_0) + \Lambda(r) - r \sum_a a \hat{f}(a) \quad (33)$$

where

$$\Lambda(r) = \log E_\infty \exp(rX) \quad (34)$$

$$E_\infty (\exp(rX - \Lambda(r)) X) = \max\left\{\frac{c^S}{n}, q_0\right\}. \quad (35)$$

2) *Discrepancy in Post-Change Conditions*: For the description in Subsubsection III-A3, we can interpret $Q(1, k) = kq\left(\hat{f}_{X_1^k}\right)$ as a random walk with i.i.d. steps of the log-likelihood ratio $\log \frac{f_{1^*}(X_i)}{f_0(X_i)}$. Finally, $S_k = \max_{\tau_k \leq i \leq k} \sum_{j=i}^k \log \frac{f_{1^*}(X_j)}{f_0(X_j)}$ becomes the CUSUM algorithm with restart.

VI. COMPARISON TO OTHER DETECTION SCHEMES

We compare IPT with various change detection algorithms proposed in the literature in terms of ROC and robust Lorden's criteria.

A. Composite Hypothesis Testing

For the hypothesis testing and transient change detection problems, we describe the simulation procedure and plot the empirical ROC, worst case probability of misdetection versus probability of false alarm, for

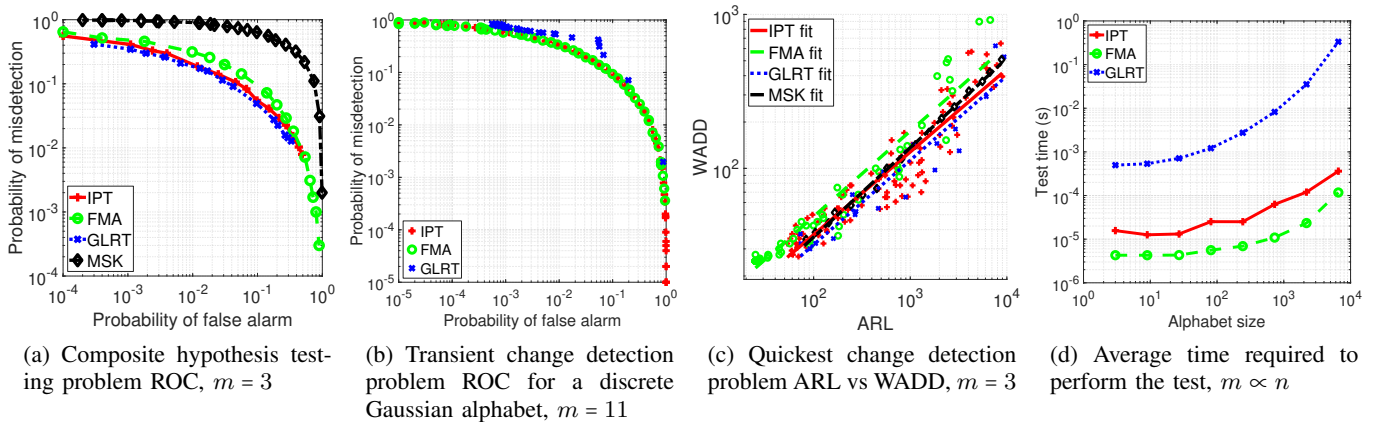


Fig. 2. IPT, FMA, GLRT and MSK are compared in terms of ROCs, robust Lorden's criteria and computational complexity.

each of the methods considered.

The ROC associated with our method is illustrated in Fig. 2a and compared with the ROC of an FMA test and GLRT. We use a ternary alphabet $\{-1, 0, 1\}$, define the pre-change distribution to be uniform and the change in mean to satisfy $\underline{q} = 0.25$. The deviations are chosen to have $n = 25$. At this stage, we determine the I-projection of f_0 on \mathcal{P}_1 . To find the optimum performance of our algorithm, we change c^D over $[2^{-8}, 2^{-3}]$. The curves in Fig. 2a show that, our proposed method is able to decrease the false positive rate compared to FMA and MSK for the same misdetection rate and performs similar to the GLRT. The area under the curves (AOC) for the worst misdetection probability are 0.0224, 0.0343, 0.0184 and 0.2915 for IPT, FMA, GLRT and MSK respectively. So, IPT reduces the AOC by 35% compared to FMA and 92% compared to MSK and has a AOC 22% larger than that of GLRT.

B. Transient Change Detection

We consider the transient change detection problem under a family of discrete Gaussian distributions. Discrete Gaussian sampling is of particular interest in lattice-based cryptography [41] and thereby in quantum-resilient cryptography [42]. For hardware based solutions proposed to generate fast and true random numbers [43], detecting a change in the distribution of generated discrete Gaussian samples could increase the reliability against failures. Here, we only consider a finite alphabet discrete Gaussian with $\mathcal{A} = \{-5, \dots, 5\}$, $m = 11$, $n = 80$, $n_\alpha = 200$ and $f(k) = N \exp\left(-\frac{k^2}{2a^2}\right)$. The known pre-change distribution has $\sigma_0^2 = 1$ and the post-change distribution satisfies $\sigma_1^2 \geq \underline{\sigma}^2 = 2$. For $q(f) = \sigma_f^2$, the set of post-change distributions $\mathcal{P}_1 = q^{-1}([\underline{\sigma}^2, \infty))$ is a convex subset of \mathcal{P} , thus we can utilize the IPT. For each method, we use a rolling window of size $w = 20 < \frac{n+1}{2}$. The worst change point t_1 is chosen among $t_1 = 1, \dots, w$ and

we randomly specify 100 different post-change distributions from \mathcal{P}_1 . Each data point in Fig. 2b is the empirical average of 10^5 worst case false alarms and worst case misdetection tests. With $\frac{20}{11} \sim 2$ samples per letter, GLRT overfits the empirical data while estimating σ_1^2 and performs worse than IPT and FMA.

C. Quickest Change Detection

For the quickest change detection problem, we test the IPT as described in Subsection V-C in terms of WADD versus ARL under the test procedure described. Then, the empirical delay curves of different algorithms are compared.

In Fig. 2c, we consider the quickest change detection performance of IPT, FMA, GLRT and MSK. The minimum WADD is plotted against the ARL under the same setting as in Subsection VI-A. IPT provides a tradeoff between FMA, MSK and GLRT with slightly worse detection delay compared to the latter. For $\gamma \geq 6000$, the minimum WADD are 335, 473, 302 and 385 for IPT, FMA, GLRT and MSK respectively. But in a scenario where computational costs would not allow a user to choose only according to the delay characteristics IPT may outperform both FMA and GLRT. A fast stream of data tracked for change at a central unit is such an example.

VII. COMPLEXITY

In this section we describe the complexity of the change detection algorithms: IPT, FMA, GLRT and MSK. We compare the computational complexity as a function of sample size n , alphabet size m and error $\epsilon \sim \|\tilde{f}^* - f^*\|_1$ involved in I-projection for IPT and maximum likelihood estimation during GLRT. FMA does not require an initialization where MSK is initialized by generating a pool of reference blocks of size n . Using a rolling window, both FMA and MSK have $\Theta(n)$ complexity⁵ to decide on the alarm time during execution. The GLRT forms a likelihood ratio for each new sample and therefore solves the following (nonconvex) optimization problem:

$$\inf_{f_1 \in \mathcal{P}_1} I(\hat{f}_{X_1^n} \| f_1) \quad (36)$$

which (if \mathcal{P}_1 is convex) is computationally equivalent to binary searching two Lagrange multipliers using an equation of m rational terms, thus it has $\mathcal{O}\left(n + m \log \frac{1}{\epsilon} \log \log \frac{1}{\epsilon}\right)$ complexity. IPT omits this projection step by initially finding the most likely deviations and then comparing any realization to this distribution.

⁵ $\mathcal{O}(g)$ is the set of functions asymptotically upper bounded by g [44], i.e. $\mathcal{O}(g) = \{f | \exists d, u_0 > 0 \text{ such that } \forall u \geq u_0, 0 \leq f(u) \leq dg(u)\}$ where, if u is a vector variable the inequalities are componentwise. Similarly, $\Theta(g) = \{f | f \in \mathcal{O}(g) \text{ and } g \in \mathcal{O}(f)\}$.

Therefore, it has only $\Theta(n+m)$ complexity. Since GLRT does not allow a recursive update, its complexity is the same for a sliding window method. Assuming the function $q(\hat{f})$ admits a recursive update in $\Theta(1)$, IPT has overall $\Theta(1)$ complexity for sliding windows. This is because $I(\hat{f}\|f^*)$ has a simple update rule. Similarly, FMA and MSK with rolling windows have $\Theta(1)$ complexity. Thus, whereas IPT, FMA and MSK have linear complexity for rolling windows and constant complexity for sliding windows, GLRT has superlinear complexity in both cases.

For a data stream with packet rate R satisfying $\frac{1}{R} \gg n + m$, GLRT can not respond to changes in the time series as fast as the other algorithms. This is especially true if a worst case detection delay is used as a cost. With a rapidly increasing demand for social media content, video, distributed computation there will be a growing need for fast heuristics than for offline exact computation.

TABLE I
TIME COMPLEXITY SCALING

	IPT	FMA	GLRT	MSK
init.	$\mathcal{O}(n + m \log \epsilon^{-1})$	$\Theta(1)$	$\Theta(1)$	$\Theta(n \log m)$
rolling	$\Theta(n + m)$	$\Theta(n)$	$\mathcal{O}(n + m \log m \log \epsilon^{-1})$	$\Theta(n)$
sliding	$\Theta(1)$	$\Theta(1)$	$\mathcal{O}(n + m \log m \log(n\epsilon)^{-1})$	$\Theta(1)$

1) *Computational Complexity:* Using Ohio Supercomputer Center Owens cluster's single node of 128GB memory and 2.40GHz clock the test times varied as in Fig. 2d [45]. We increased the alphabet and sample sizes proportionally and averaged over the total number of tests. IPT performed 25 to 670 times faster than the GLRT where the gain from test time increased with the alphabet size.

VIII. APPLICATIONS TO ECONOMICS AND CLIMATE DATA

In this section, we apply our outlier detection scheme, based on checking the RLLF of the most likely outlier distribution to the empirical distribution of the given period. We would like to test whether shifts of historic value are merely outliers or it happens due to an exogenous factor. To obtain the ground truth, we will use the long-term empirical distribution. Thus, the underlying assumption we make in the following data analysis is that, for r shifts of length n the overall cumulative duration of deviations is much smaller compared to the size of the data: $rn \ll T$. To analyze a period of interest, we pick a duration n and a threshold, c^S , that matches the mean of the underlying data over that period.

A. Daily Returns of Portfolios by Size

We focus on portfolio returns with different market caps [46] to identify segments with a worse performance compared to the average behavior, sustained over a substantial duration of time. The primary question we answer is the following: “is the identified negative sequence a rare event generated by the statistical nature of the market or is it driven by some exogenous factors with the potential to cause a financial crisis?” The monthly averages of the return vary between 0.8 – 0.9. We quantized the daily percentage portfolio returns for different market caps over July 1926-March 2019 and selected a threshold c^S below the average returns. We then computed the I-projection of the long term empirical distribution over the distributions with mean less than the threshold. The $n = 6$ month moving average and the KL divergence of the observation window against the I-projection is given in Fig. 3. With the right choice of n and c we were able to identify historically significant events like the great depression or the 2009 financial crisis. One of the most significant finding is that, even though there are other periods with mean return as low as these periods, the RLLF clearly differentiated these periods from the crises times, showing strength of our scheme.

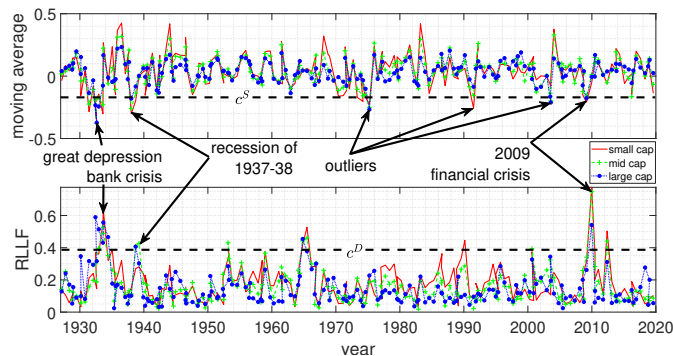


Fig. 3. The moving average and the RLLF for the market indices with different market capitalization

B. Oil Industry Market Index

Similar to Subsection VIII-A, this time we focus on market index for oil industry. We use monthly return series of Oil Industry Portfolio Index, obtained from Kenneth French Data Library [46] and covers the period from 1926-07-01 to 2019-03-31. The historical monthly return of the index is 0.89. We have run the experiment with $c^S = -0.1\%$ and $n = 12$ months. In Fig. 4, we have 7 first threshold crosses, only 2 of which classified as change. The greatest divergence is for the 2009 oil crisis period. Unfortunately, the 1973 oil embargo is disregarded as a false alarm. This example shows the requirement of tuning for our method.




Fig. 4. The moving average and the RLLF for the oil industry market index

C. Oxygen Isotope Data

We analyze climate data collected from a polar cap that spans ~ 1800 years to gain insight of our method and historic climate data [47]. Fig. 5 uses a moving average window of $n = 20$ years of quantized data of oxygen isotope density over the years 226 – 2009. We then select thresholds c^S , half a standard deviation above the mean, and c^D and determine the I-projections. Note that the only point IPT classifies as change is the last 20-year period of above norm d18O levels. Although the threshold selections are not unique this intuitive approach with these hyperparameters gives insight that the last 20-year period is less likely to be an outlier rather than an effect of an exogenous change, i.e. man made.




Fig. 5. The moving average and RLLF for the oxygen isotopes and annual ice accumulation

IX. CONCLUSION

In a variety of applications, it is important to identify shifts from the typical behavior. However, not every shift from the norm marks a regime change or a deviation; instead it could be an outlier. In this paper, we developed a method that differentiates between outliers and change. To achieve that, we used results from large deviations theory, which gives us statistical characterization for the outliers. Our method tests a given string against the most likely way an outlier occurs and marks deviations using detection-theoretic tools.

Our method uses control parameters c^S and c^D to shape the decision region to determine the outliers. We have bounded the probability of false positive and misdetection as functions of these control parameters and shown that the detection delay is asymptotically optimal up to a constant as ARL goes to ∞ .

We also emphasized IPT's computational simplicity in that it reduces GLRT's requirement to run non-convex and non-recursive optimization problems to solving convex optimization problems during initialization. This has proven to reduce the average running time per test 25 to 670 fold for varying alphabet sizes.

We also applied our algorithm to a variety of applications and drew insights from the observed time series data. Our method verified the global warming phenomenon as highly unlikely to be an outlier, giving credence for it to be caused by factors (such as those man-made) leading to a statistical change in the indicator variables. Similarly, we tried to extract historically significant exogenous events for the market within the data.

We have omitted proves that complicate the exposition whenever necessary but we believe similar results can be extended to the continuous alphabet case via binning.

Future work will be on proposing a metric that naturally and objectively combines sample delay for detection with computational complexity to address true delays as observed in high traffic data centers that detect changes in popularity, inappropriate content, fake accounts or attacks.

REFERENCES

- [1] 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, June 2019.
- [2] J. Chen, "Testing for a change point in linear regression models," *Communications in Statistics - Theory and Methods*, vol. 27, no. 10, pp. 2481–2493, 1998. [Online]. Available: <https://doi.org/10.1080/03610929808832238>
- [3] 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, May 2018.
- [4] I. Nikiforov, V. Varavva, and V. Kireichikov, "Application of statistical fault detection algorithms to navigation systems monitoring," *Automatica*, vol. 29, no. 5, pp. 1275 – 1290, 1993.
- [5] P. Perera and V. M. Patel, "Efficient and low latency detection of intruders in mobile active authentication," *IEEE Transactions on Information Forensics and Security*, vol. 13, no. 6, pp. 1392–1405, June 2018.
- [6] —, "Quickest intrusion detection in mobile active user authentication," in *2016 IEEE 8th International Conference on Biometrics Theory, Applications and Systems (BTAS)*, 2016, pp. 1–8.
- [7] G. Baldewijns, S. Luca, W. Nagels, B. Vanrumste, and T. Croonenborghs, "Automatic detection of health changes using statistical process control techniques on measured transfer times of elderly," in *2015 37th Annual International Conference of the IEEE Engineering in Medicine and Biology Society (EMBC)*, Aug 2015, pp. 5046–5049.

- [8] V. Nika, P. Babyn, and H. Zhu, “Change detection of medical images using dictionary learning techniques and principal component analysis,” *Journal of medical imaging (Bellingham, Wash.)*, vol. 1, no. 2, pp. 024 502–024 502, Jul 2014.
- [9] H. V. Poor and O. Hadjiliadis, *Quickest Detection*. Cambridge University Press, 2008.
- [10] M. Basseville and I. V. Nikiforov, *Detection of Abrupt Changes: Theory and Application*. Upper Saddle River, NJ, USA: Prentice-Hall, Inc., 1993.
- [11] 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, Feb 2019.
- [12] A. G. Tartakovsky, “Asymptotically optimal quickest change detection in multistream data—part 1: General stochastic models,” *Methodology and Computing in Applied Probability*, vol. 21, no. 4, pp. 1303–1336, Dec 2019.
- [13] 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, March 2011.
- [14] T. S. Lau, W. P. Tay, and V. V. Veeravalli, “Quickest change detection with unknown post-change distribution,” in *2017 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, March 2017, pp. 3924–3928.
- [15] G. Lorden and M. Pollak, “Sequential change-point detection procedures that are nearly optimal and computationally simple,” *Sequential Analysis*, vol. 27, no. 4, pp. 476–512, 2008.
- [16] S. Li, Y. Xie, H. Dai, and L. Song, “M-statistic for kernel change-point detection,” in *Advances in Neural Information Processing Systems*, C. Cortes, N. Lawrence, D. Lee, M. Sugiyama, and R. Garnett, Eds., vol. 28. Curran Associates, Inc., 2015.
- [17] I. Nikiforov, “Quickest multidecision abrupt change detection with some applications to network monitoring,” in *Distributed Computer and Communication Networks*, V. Vishnevsky and D. Kozyrev, Eds. Cham: Springer International Publishing, 2016, pp. 94–101.
- [18] L. Xie, G. V. Moustakides, and Y. Xie, “First-order optimal sequential subspace change-point detection,” in *2018 IEEE Global Conference on Signal and Information Processing (GlobalSIP)*, 2018, pp. 111–115.
- [19] V. V. Veeravalli and T. Banerjee, “Chapter 6 - quickest change detection,” in *Academic Press Library in Signal Processing: Volume 3*, ser. Academic Press Library in Signal Processing, A. M. Zoubir, M. Viberg, R. Chellappa, and S. Theodoridis, Eds. Elsevier, 2014, vol. 3, pp. 209–255.
- [20] O. Hadjiliadis and V. Moustakides, “Optimal and asymptotically optimal cusum rules for change point detection in the brownian motion model with multiple alternatives,” *Theory of Probability & Its Applications*, vol. 50, no. 1, pp. 75–85, 2006.
- [21] T. Banerjee and V. V. Veeravalli, “Data-efficient minimax quickest change detection with composite post-change distribution,” *IEEE Transactions on Information Theory*, vol. 61, no. 9, pp. 5172–5184, 2015.
- [22] G. V. Moustakides, “Multiple optimality properties of the shewhart test,” *Sequential Analysis*, vol. 33, no. 3, pp. 318–344, 2014.
- [23] T. S. Lau and W. P. Tay, “Quickest change detection under a nuisance change,” in *2018 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, April 2018, pp. 6643–6647.
- [24] 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, March 2018.
- [25] S. Skaperas, L. Mamatas, and A. Chorti, “Early video content popularity detection with change point analysis,” in *2018 IEEE Global Communications Conference (GLOBECOM)*, Dec 2018, pp. 1–7.
- [26] 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, June 2019.
- [27] 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, 2020.

- [28] D. Sargun and C. E. Koksak, "Separating an outlier from a change," in *2019 53rd Asilomar Conference on Signals, Systems, and Computers*, Nov 2019.
- [29] G. Lorden, "Procedures for reacting to a change in distribution," *Ann. Math. Statist.*, vol. 42, no. 6, pp. 1897–1908, 12 1971.
- [30] H. V. Poor, *An Introduction to Signal Detection and Estimation (2nd Ed.)*. Berlin, Heidelberg: Springer-Verlag, 1994.
- [31] R. G. Gallager, *Stochastic processes: theory for applications*. Cambridge University Press, 2016.
- [32] A. Dembo and O. Zeitouni, *Large Deviations Techniques and Applications*. Springer Berlin Heidelberg, 2010.
- [33] A. Sen and M. S. Srivastava, "On tests for detecting change in mean," *The Annals of Statistics*, vol. 3, no. 1, pp. 98 – 108, 1975.
- [34] S. Jewell, P. Fearnhead, and D. Witten, "Testing for a change in mean after changepoint detection," 2020.
- [35] T. Banerjee and V. V. Veeravalli, "Data-efficient quickest change detection with unknown post-change distribution," in *2014 IEEE International Symposium on Information Theory*, 2014, pp. 741–745.
- [36] A. G. Tartakovsky and A. S. Polunchenko, "Quickest changepoint detection in distributed multisensor systems under unknown parameters," in *2008 11th International Conference on Information Fusion*, 2008, pp. 1–8.
- [37] A. Beck and M. Teboulle, "Mirror descent and nonlinear projected subgradient methods for convex optimization," *Operations Research Letters*, vol. 31, no. 3, pp. 167–175, 2003.
- [38] D. P. Bertsekas, "Projected newton methods for optimization problems with simple constraints," *SIAM Journal on Control and Optimization*, vol. 20, no. 2, pp. 221–246, 1982.
- [39] S. Boyd and L. Vandenberghe, *Convex Optimization*. Cambridge University Press, 2004.
- [40] T. Bojdecki, "Probability maximizing approach to optimal stopping and its application to a disorder problem," *Stochastics*, vol. 3, no. 1-4, pp. 61–71, 1980.
- [41] A. Karmakar, S. S. Roy, O. Reparaz, F. Vercauteren, and I. Verbauwhede, "Constant-time discrete gaussian sampling," *IEEE Transactions on Computers*, vol. 67, no. 11, pp. 1561–1571, Nov 2018.
- [42] J. Howe, A. Khalid, C. Rafferty, F. Regazzoni, and M. O'Neill, "On practical discrete gaussian samplers for lattice-based cryptography," *IEEE Transactions on Computers*, vol. 67, no. 3, pp. 322–334, March 2018.
- [43] Y. Liu, X. Li, R. C. C. Cheung, S. Chan, and H. Wong, "High-speed discrete gaussian sampler with heterodyne chaotic laser inputs," *IEEE Transactions on Circuits and Systems II: Express Briefs*, vol. 65, no. 6, pp. 794–798, June 2018.
- [44] T. H. Cormen, C. E. Leiserson, R. L. Rivest, and C. Stein, *Introduction to Algorithms, Third Edition*, 3rd ed. The MIT Press, 2009.
- [45] O. S. Center, "Ohio supercomputer center," 1987. [Online]. Available: <http://osc.edu/ark:/19495/f5s1ph73>
- [46] K. R. French, "Data library," Tuck School of Business, Dartmouth University, Tech. Rep., 2019. [Online]. Available: <http://mba.tuck.dartmouth.edu/pages/faculty/ken.french/index.html>
- [47] L. Thompson, E. Mosley Thompson, M. Davis, V. Zagorodnov, I. Howat, V. Mikhalenko, and P.-N. Lin, "Annually resolved ice core records of tropical climate variability over the past 1800 years," *Science*, vol. 340(6135), pp. 945–950, 2013.