

Bandit Change-Point Detection for Real-Time Monitoring High-Dimensional Data Under Sampling Control

Wanrong Zhang
Harvard University
and
Yajun Mei
Georgia Institute of Technology

Abstract

In many real-world problems of real-time monitoring high-dimensional streaming data, one wants to detect an undesired event or change quickly once it occurs, but under the sampling control constraint in the sense that one might be able to only observe or use selected components data for decision-making per time step in the resource-constrained environments. In this paper, we propose to incorporate multi-armed bandit approaches into sequential change-point detection to develop an efficient bandit change-point detection algorithm based on the limiting Bayesian approach to incorporate a prior knowledge of potential changes. Our proposed algorithm, termed Thompson-Sampling-Shiryaev-Roberts-Pollak (TSSRP), consists of two policies per time step: the adaptive sampling policy applies the Thompson Sampling algorithm to balance between exploration for acquiring long-term knowledge and exploitation for immediate reward gain, and the statistical decision policy fuses the local Shiryaev-Roberts-Pollak statistics to determine whether to raise a global alarm by sum shrinkage techniques. Extensive numerical simulations and case studies demonstrate the statistical and computational efficiency of our proposed TSSRP algorithm.

Keywords: adaptive sampling; change-point detection; partially observed variables; Shiryaev-Roberts procedure; Thompson sampling

1 Introduction

Real-time monitoring high-dimensional streaming data under sampling control constraints appears in many important applications such as intrusion detection in computer networks (Bass, 1999), event detection in social networks (Viswanath et al., 2014), epidemic disease outbreak monitoring (Yang et al., 2015), anomaly detection in manufacture processes (Ding et al., 2006). In these applications, one often can only observe or use selected components of the data for decision-making due to the capacity limitation in data acquisition, transmission, processing, or storage. For instance, the sensor devices might have limited battery powers; thus, one might want to use a subset of sensors per time step over a long period instead of using full sensors simultaneously over a short period. Likewise, while sensing is usually cheap, the communication bandwidth is often limited from remote sensors to the fusion center that makes a global decision. The fusion center might prioritize certain local sensors to send local information for decision making. Also, in many applications such as quality engineering or biosurveillance, one faces the design issue and needs to decide which variables or patients to be measured to detect the defect or disease outbreak more efficiently.

In this work, we investigate how to efficiently real-time monitor high-dimensional streaming data under resource constraints. We assume that the full data from a system is a K -dimensional random vector $\mathbf{X}_t = (X_{1,t}, \dots, X_{K,t})$ at each time step, but we can only observe q out of K components per time step. Here the component $X_{k,t}$, with $k = 1, \dots, K$ and $t = 1, 2, \dots$, can be either the raw data from local sensors or the derived features such as wavelet coefficients, principal components. Initially, the system is in control in the sense

that $X_{k,t}$ follows a probability density function f_k . At some unknown time ν , an event may occur and change the distributions of a sparse subset of the components. Our goal is to design an efficient algorithm to adaptively decide which variable to sample at each time step, and when to raise a global alarm to indicate the possible occurrence of the change.

Without resource constraints, monitoring fully observed streaming data has raised much attention in the statistical quality control (SPC) and sequential change-point detection literature, see Zou and Qiu (2009); Li (2019, 2020); Li et al. (2020); Zou et al. (2015). The existing work generally falls into two frameworks: the cumulative sum (CUSUM) type method, which is based on the generalized likelihood ratio (GLR) framework; and the Shiryaev-Roberts type method, which is based on the Bayesian framework. For classical research on one-dimensional data streams, see Shiryaev (1963); Lorden (1971); Pollak (1985); Lai (1995, 1998); Basseville and Nikiforov (1993); Poor and Hadjiladis (2008); Tartakovsky et al. (2014). For recent research on high-dimensional data streams with fully observed data, see Zhang and Siegmund (2012); Xie and Siegmund (2013); Wang and Mei (2015); Cho and Fryzlewicz (2015); Chan (2017); Chu and Chen (2019). Additionally, another framework is to monitor each data stream separately by computing respective local detection statistics and then fuse local statistics into a global-level monitoring statistic, see Mei (2010, 2011); Liu et al. (2019); Li (2020). This framework can balance the tradeoff between computational efficiency and statistical efficiency.

Real-time monitoring high-dimensional *partially* observed data streams under the sampling control has been studied in the literature of statistical process control in applied statistics. A prominent line of work is based on the CUSUM procedure for observed local

streams with an artificially introduced compensation parameter for the unobserved local stream, see Liu et al. (2015); Xian et al. (2018); Wang et al. (2018); Xian et al. (2021). While the compensation parameter can increase the chance of exploring unobserved local streams, tuning the parameter is challenging. Another line of work leverages extra information such as the correlation structure to approximate unobserved local streams and then to plug in the standard monitoring methods of fully observed data, see Zhang and Hoi (2019); Nabhan et al. (2021).

We propose a bandit change-point detection algorithm for efficient real-time monitoring of high-dimensional streaming data under the sampling control. Our contributions are twofold: (i) We incorporate prior knowledge of potential changes to update unobserved local streams by treating the likelihood ratios of unobserved data as one. (ii) We incorporate the Thompson Sampling algorithm in the multi-armed bandit (MAB) problem into the Shiryaev-Roberts-Pollak procedure in the sequential change-point detection literature. While Bayesian methods often involve extensive computations, our method, termed as Thompson-Sampling-Shiryaev-Roberts-Pollak (TSSRP) algorithm, is computationally efficient when monitoring high-dimensional data when the data streams are mutually independent and we have some prior knowledge on the post-change distribution. The limiting Bayesian framework allows our algorithm to have a natural interpretation and avoid the tuning of artificial tuning parameters such as the compensation parameters for unobserved data. It can balance the tradeoff between exploiting the observed local components that maximize the immediate detection performance and exploring not-been-monitored local components that might provide new information to improve future detection performance.

In particular, our proposed TSSRP algorithm performs similar to random sampling in the in-control state when no changes occurring, but becomes a greedy sampling on those affected local components in the out-of-control state when a change occurs. Numerical simulations and case studies show the efficiency of our proposed TSSRP algorithm.

The classical MAB problem focuses on developing algorithms to balance the tradeoff between exploration for acquiring long-term knowledge and exploitation for immediate reward gain, see Lai and Robbins (1985); Robbins (1985); Scott (2010); Gittins et al. (2011); Bubeck and Cesa-Bianchi (2012); Agrawal and Goyal (2012); Cao et al. (2019); Zhao (2019) and references therein. Our work and the classical MAB problem both deal with the dynamical/adaptive sampling strategy that samples those local streams with the largest values of some suitable local statistics. Nevertheless, our work is different from the existing research on the MAB with non-stationary or piecewise constant rewarding functions, see Cao et al. (2019); Ghatak (2020), because our primary objective is to minimize the average detection delay subject to controlling false alarm rate, whereas the bandit problems minimize cumulative regret. In summary, we apply the bandit ideas to develop a new sequential change-point detection algorithm for monitoring partially observed data.

The remainder of this paper is organized as follows. In Section 2, we provide the mathematical formulation of our problem and also review the background of the multi-armed bandit problem and the sequential change-point detection problem. Next, we introduce our proposed method and develop its theoretical properties in Section 3. Then we evaluate the performance of our proposed algorithm through simulation studies and real data case studies in Section 4 and 5, respectively. Concluding remarks are included in Section 6.

We provide the detailed proofs of all theorems in the Supplementary Materials, which also include additional numerical experiments.

2 Problem Formulation and Backgrounds

In this section, we present the mathematical formulation of real-time monitoring high-dimensional streaming data in resource-constrained environments in Subsection 2.1. Then we provide a brief review of the Thompson Sampling algorithm for the multi-armed bandit problem in Subsection 2.2, followed by the review of the Bayesian approach for sequential change-point detection in Subsection 2.3.

2.1 Problem Formulation

Suppose we are monitoring K independent data streams in a system. Let $X_{k,t}$ denote the observation of the k -th data stream at time t for $t = 1, 2, \dots$ and $k = 1, 2, \dots, K$. Here each data stream $X_{k,t}$ can be the raw data itself or its derived features such as wavelet coefficients, principal components. Samples from data streams are assumed mutually independent. Each data stream generates identically distributed data from a specific distribution f_{θ_k} . At some unknown change time $\nu \in \{1, 2, \dots\}$, an undesirable event occurs and changes the distributions of some data streams abruptly in the sense of changing the values of the parameters θ_k . Conditional on the change time $\nu > 1$, for those affected data streams, the observation $X_{k,1}, \dots, X_{k,\nu-1}$ are independent and identically distributed (iid) with density $f_{\theta_{k,0}}$ while $X_{k,\nu}, X_{k,\nu+1}, \dots$ are iid with another density $f_{\theta_{k,1}}$, where $\theta_{k,1} > \theta_{k,0}$, and the case where $\theta_{k,1} < \theta_{k,0}$ can be handled similarly. For those unaffected data streams, all

observations $X_{k,1}, X_{k,2} \dots$ are iid with density $f_{\theta_{k,0}}$. Here we do not know which subset of the local streams changed, and we assume that the affected local streams are sparse. In practice, practitioners usually specify $f_{\theta_{k,1}}$ as the interested-smallest magnitude of a change to be detected.

Under sampling control, we can only observe or use selected partial data for decision making. We assume that only $q < K$ data streams can be selected to collect data at each time t . Mathematically, let $\delta_{k,t}$ be the indicator function that $\delta_{k,t} = 1$ if and only if the k -th data stream $X_{k,t}$ is selected at time step t . The resource constraint implies that $\sum_{k=1}^K \delta_{k,t} = q$ at each time step $t = 1, 2, \dots$. Imagine that we put q sensors onto K locations, then the $\delta_{k,t}$ can be thought as whether to put a sensor to the k -th data stream at time t . Let S_t denote the locations where $\delta_{k,t} = 1$, and we refer to it as the sensor layouts. Under our notation, the observed data can be represented as $\{X_{k,t}^*\} = \{X_{k,t}\delta_{k,t}\}$, $k = 1, 2, \dots, K$.

For the change-point detection problem under sampling control, a statistical scheme consists of two policies per time step. One is the adaptive sampling policy δ that decides the observable location $\delta_{k,t}$, and the other is the statistical decision policy, often defined as a stopping time T , that raises an alarm based on the observations available $\{X_{k,t}^*\} = \{X_{k,t}\delta_{k,t}\}_{1 \leq k \leq K, 1 \leq t \leq T}$. Our objective is to design a scalable and efficient statistical scheme of (δ, T) that minimizes the average detection delay

$$D(T) = \sup_{1 \leq \nu < \infty} \mathbb{E}_\nu(T - \nu | T \geq \nu), \quad (1)$$

subject to the Average Run length (ARL) to False Alarm constraint

$$\mathbb{E}(T|\nu = \infty) \geq \gamma, \tag{2}$$

where γ is a pre-specified constant.

It is worth noting that our problem connects to the multi-armed bandit problem in the sampling policy; however, they are fundamentally different due to different performance criteria.

2.2 Thompson Sampling for Multi-Armed Bandit

In this subsection, we briefly review the multi-armed bandit (MAB) problem, first introduced by Robbins (1952), and one of the most popular algorithms, Thompson Sampling (Thompson, 1933). Under a classical setting of MAB, a gambler can play one of K slot machines (or arms) for $K \geq 2$, but she or he has no prior knowledge about which machine has a potentially higher reward. The only way to learn rewards is to play the machines. The problem of interest is how the gambler decides which arm to play at each time step, to maximize the total rewards through N plays.

One of the most popular MAB algorithms is Thompson Sampling, which is a natural Bayesian algorithm. It has been widely used for personalized advertisements and product recommendation (Agrawal and Goyal, 2013), as its efficiency has been well demonstrated in many real-world applications, especially in the high-dimensional setting. See Scott (2010), Chapelle and Li (2011). In particular, Agrawal and Goyal (2012) showed that the Thompson Sampling algorithm asymptotically minimizes the expected regret.

The idea of Thompson Sampling is to sample arms based on the largest values of the random realizations of the posterior distributions instead of the posterior means. Specifically, at each time step, after updating the posterior distribution of the mean θ_k for each arm, we randomly sample a realization from the posterior distributions, denoted by $\hat{\theta}_k$ from the k -th arm. Then we select the arm with the largest random realization, i.e., $\arg \max_{1 \leq k \leq K} \hat{\theta}_k$. This allows us to have better chances to sample those arms with fewer observations, thereby balancing the tradeoff between the exploration for acquiring long-term knowledge and the exploitation for immediate reward gain.

In the multi-armed bandit problem, when we are allowed to observe q arms each time, it is natural to extend the original Thompson Sampling algorithm to select the q -largest realizations. Such an approach often holds nice properties under reasonable conditions, see Anantharam et al. (1987); Pandelis and Teneketzis (1999); Kaufmann et al. (2016). Thus we will adopt the Thompson sampling with q -largest realizations in our context.

2.3 Shiryaev-Roberts Procedure

We now review the Bayesian approach for the simplest sequential change-point detection problem pioneered by Shiryaev (1963), as well as the corresponding limiting Bayesian approach. See Roberts (1966); Pollak (1985, 1987). Consider the simplest univariate case when we observe a sequence of independent observations X_1, X_2, \dots , whose distribution might change from f_{θ_0} to f_{θ_1} at some unknown time ν . Since the goal is to detect the change time ν quickly, the statistical procedure is defined as a stopping time T with respect to the observed data $\{X_t\}_{t \geq 1}$, where $\{T = t\}$ means that we raise the alarm at time

t to indicate that a change has occurred up to time t .

Under the Bayesian formulation, it is assumed that the change-point ν has a geometric prior distribution:

$$P(\nu = t) = p(1 - p)^{t-1} \quad \text{for } t = 1, 2, 3, \dots \quad (3)$$

where $0 < p < 1$ is a pre-specified constant. Moreover, conditional on the (unknown) change-point ν , the pre-change observations, $X_1, \dots, X_{\nu-1}$, are iid with density f_{θ_0} and are independent of the post-change observations, $X_\nu, X_{\nu+1}, \dots$ which are iid with density f_{θ_1} . Assume that the cost of per post-change observation is $c > 0$. Then the Bayesian formulation of sequential change-point detection is to find a statistical procedure T that minimizes the Bayes risk $P(T < \nu) + c\mathbb{E}(T - \nu)^+$.

Shiryaev (1963) first solved this problem, and the Bayesian solution is to raise an alarm at the first time when the posterior probability of change having occurred, i.e., $P(\nu \leq t | X_1, \dots, X_t)$, is greater than a certain threshold. Under the limiting Bayes framework, one considers the test statistic of the form $P(\nu \leq t | X_1, \dots, X_t) / p(1 - P(\nu \leq t | X_1, \dots, X_t))$ as p goes to zero. This yields the so-called Shiryaev-Roberts procedure (Roberts (1966)) that raises an alarm at time

$$T_A = \inf\{t | R_t \geq A\}, \quad (4)$$

where R_t is the Shiryaev-Roberts statistic defined as

$$R_t = \sum_{j=1}^t \prod_{i=j}^t \frac{f_{\theta_1}(X_i)}{f_{\theta_0}(X_i)}, \quad (5)$$

and the threshold A is a pre-specified constant. Pollak (1985, 1987) showed that this procedure enjoys nice asymptotic minimax properties, i.e., minimize the worst average

detection delay in (1) up to within an $o(1)$ term subject to the ARL to false alarm constraint in (2), as γ goes to ∞ .

3 Bandit Change-Point Detection

In this section, we present our proposed TSSRP algorithm for the real-time monitoring high-dimensional streaming data under sampling control. Our proposed algorithm can be thought of as the limit of Bayesian procedures that adapt the Thompson sampling policy of sampling local streams based on the random realizations of the posterior distributions.

We assume the data contains K independent data streams. The local change time ν_k of the k -th local data stream has a prior Geometric(p) distribution. The k -th local stream has an initial prior probability $\Pi_{k,0}$ of change and $\Pi_{k,0}$'s are mutually independent, and identically distributed from a common prior $G = G_p$. We can extend the idea to other non-homogeneous scenarios, e.g., in quality control of K -stages manufacturing process where some stages are more prone to defect. Specifically, the distribution of the local change time ν_k is as follows:

$$P(\nu_k = 0) = \Pi_{k,0}, \quad P(\nu_k = t) = (1 - \Pi_{k,0})p(1 - p)^{t-1} \quad \text{for } t = 1, 2, 3, \dots, \quad (6)$$

where $\Pi_{k,0}$ can be either a constant, e.g., $\Pi_{k,0} \equiv 0$, or a random variable that has a distribution G_p .

After taking observations at the time step t , we update the posterior distribution of ν_k , denoted by $\Pi_{k,t} = P(\nu_k \leq t | \text{Observed Data})$. This computation is straightforward since the raw data $X_{k,t}$ is distributed as $f_{\theta_{k,0}}I(\nu_k < t) + f_{\theta_{k,1}}I(\nu_k \geq t)$, although it is observable if and

only if the sampling indicator $\delta_{k,t} = 1$. Next, we combine the local posterior distributions $\Pi_{k,t}$'s together to decide if we would raise a global alarm. If yes, then we stop taking any observations. If no, then we will proceed to the next time step — as in the Thompson sampling algorithm, we decide to observe those local streams in time step $t + 1$ with the largest values of the posterior values $\Pi_{k,t}$'s if the initial values $\Pi_{k,0}$ are constant, or their random realizations if the initial values $\Pi_{k,0}$ are random variables.

Implementing the Bayesian algorithm in a naive way is computationally infeasible. Our algorithm overcomes this challenge by leveraging the property that the limit of the Bayesian algorithms as p goes to 0 has a mathematical equivalent representation that is computationally scalable. Therefore, our algorithm is both statistically and computationally efficient.

We present our proposed TSSRP methodology in Subsection 3.1 and discuss the choice of parameters and prior distribution in Subsection 3.2. We develop the theoretical properties of our proposed TSSRP algorithm including its connection to the Bayesian procedures in Subsection 3.3.

3.1 Methodology Development

In the context of real-time monitoring high-dimensional streaming data under sampling control, a statistical procedure consists of two policies per time step: (1) the adaptive sampling policy to decide the observation location; (2) the statistical decision policy to raise a global alarm based on the observed data. A common challenge in both components or policies is how to construct local statistics for each local stream that can guide us to make efficient decisions for both adaptive sampling and statistical decision policies. Our

proposed method’s key novelty is to recursively update two-dimensional local statistics over time that allow conveniently implement the Thompson Sampling.

3.1.1 Local Statistics

We propose to recursively compute two-dimensional local statistics, denoted by $R_{k,t}$ and $L_{k,t}$, at the k -th local data steam at each time step $t = 1, 2, \dots$, where $R_{k,t}$ mimics the classical Shiryaev-Robert statistics

$$R_{k,t} = \begin{cases} (R_{k,t-1} + 1) \frac{f_{\theta_{k,1}}(X_{k,t})}{f_{\theta_{k,0}}(X_{k,t})}, & \text{if } \delta_{k,t} = 1; \\ R_{k,t-1} + 1, & \text{if } \delta_{k,t} = 0. \end{cases} \quad (7)$$

with initial value $R_{k,t=0} = 0$, and the statistics $L_{k,t}$ mimics the likelihood ratio function and

$$L_{k,t} = \begin{cases} L_{k,t-1} \frac{f_{\theta_{k,1}}(X_{k,t})}{f_{\theta_{k,0}}(X_{k,t})}, & \text{if } \delta_{k,t} = 1; \\ L_{k,t-1}, & \text{if } \delta_{k,t} = 0. \end{cases} \quad (8)$$

with initial value $L_{k,t=0} = 1$.

At a high-level, the local statistic $R_{k,t}$ mainly provides the evidence how likely a local change has occurred, whereas the other local statistic $L_{k,t}$ is related to the number of samples taken at a given local data stream. If $\delta_{k,t} = 1$, i.e., if one takes observations from that specific local data streams, then the update on $R_{k,t}$ and $L_{k,t}$ follows the classical Shiryaev-Robert or likelihood statistics, respectively. On the other hand, if $\delta_{k,t} = 0$, i.e., if we do not take local observations, then we recursively update $R_{k,t}$ and $L_{k,t}$ by adding or multiplying the constant 1, respectively. The intuition is to treat $\frac{f_{\theta_{k,1}}(X_{k,t})}{f_{\theta_{k,0}}(X_{k,t})}$ as 1 if $X_{k,t}$ is missing.

Note that the definition or computation of $(R_{k,t}, L_{k,t})$ depends on which local sensors will be observed, i.e., the values of sampling indicator variables $\delta_{k,t}$'s, which will be defined in the next subsection.

3.1.2 Adaptive Sampling Policy

Our proposed adaptive sampling policy is as follows.

At each time step $t = 0, 1, 2, \dots$, we compute the two-dimensional local statistics, $(R_{k,t}, L_{k,t})$, based on the observed data streams, and we also sample a randomized value $\tilde{R}_{k,t}$ from a pre-specified prior distribution G . The $\tilde{R}_{k,t}$ can be treated as the “initial values.”

When G is a point mass density of 0, $\tilde{R}_{k,t} \equiv 0$ for all $k = 1, \dots, K$. Otherwise, $\tilde{R}_{k,t}$ can be different across sensors, which can be viewed as the prior knowledge of how likely a local data stream is likely affected by the change. Next, we compute a real-valued local statistic that determines the sampling policies:

$$R_{k,t}^* = R_{k,t} + L_{k,t} \tilde{R}_{k,t}. \quad (9)$$

Finally, at time step $t + 1$, we follow the Thompson Sampling to adaptively choose the local data streams with the largest q values of $R_{(k),t}^*$ in (9). Let $l_{(k),t+1}$ denote the corresponding index of the k th largest values, then the new sensor layout will be $S_{t+1} = \{l_{(1),t+1}, \dots, l_{(q),t+1}\}$ at time $t + 1$.

Let us provide a high-level rationale for our proposed adaptive sampling policy. First, note that $R_{k,t}^*$ in (9) can be thought of as a randomized version of $R_{k,t}$ and allows us to balance better the tradeoff between those local streams having larger observed $R_{k,t}$ and

those local streams having fewer observations. Second, our sampling policy is computationally efficient, since it is based on the recursive updates of the two-dimensional local statistics, $(R_{k,t}, L_{k,t})$. Finally, our sampling policy is the Thompson sampling method under the limiting Bayes framework, since the larger the $R_{k,t}^*$ value, the larger the realization of the posterior distribution of a local change.

3.1.3 Global Decision

Our proposed global decision policy is to raise a global alarm based on the largest r values of the local statistics $R_{k,t}$ in (7), and is defined as the stopping time

$$T = \inf\{t \geq 1 : \sum_{k=1}^r R_{(k),t} \geq A\}, \quad (10)$$

where r is a pre-specified parameter, and A is a pre-specified constant so as to satisfy the false alarm constraint in (2). Here $R_{(1),t} \geq \dots \geq R_{(k),t} \geq \dots \geq R_{(K),t}$ denote the decreasing order of the local statistics $R_{k,t}$ in (7).

We should acknowledge that there are many other ways to raise a global decision. For instance, for local statistics in the summation, we can use the randomized version $R_{(k),t}^*$ or the logarithm version $\log R_{(k),t}$. Moreover, there are different ways to use the shrinkage transformation to combine local statistics to raise a global alarm; see Mei (2011) and Liu et al. (2019). Based on our extensive simulation experiences, the stopping time in (10) is stable and outperforms other stopping rules in most cases. The discussion on comparison with other stopping rules is deferred to the supplementary material Section 3.

Algorithm 1 Thompson-Sampling-Shiryayev-Roberts-Pollak (TSSRP) algorithm

Parameters: the number r , the number of observed sensors q , a prior distribution G and the stopping threshold A .

Input: K data streams

Initialize: Set $R_{k,t} = 0$, $L_{k,t} = 1$, and sample $\tilde{R}_{k,t}$ from G for all $k = 1, 2, \dots, K$. Randomly sample q data streams as the initial layout S_1

Algorithm: In each round $t \leftarrow 1, 2, \dots$ do the following:

- (1) based on the current sensor layout S_t , recursively update two-dimensional local statistics $(R_{k,t}, L_{k,t})$ in (7) and (8)
 - (2) For each data stream k , sample the “initial” value $\tilde{R}_{k,t}$ from G , and calculate the local sampling statistics $R_{k,t}^*$ in (9)
 - (3) Order the local sampling statistics $R_{k,t}^*$, $k = 1, 2, \dots, K$, from the largest to the smallest, and let $l_{(k),t}$ denote the variable index of the order statistics $R_{(k),t}^*$
 - (4) Update the sensor layout = $\{l_{(1),t}, \dots, l_{(q),t}\}$
 - (5) Check if the criterion of the stopping time in (10) is reached. If yes, stop and raise a global alarm. If not, proceed to the next iteration.
-

3.1.4 Summary of Proposed Algorithm

We summarize the proposed Thompson-Sampling-Shiryayev-Roberts-Pollak (TSSRP) in Algorithm 1.

Our proposed TSSRP method is not only statistical efficient as a limiting Bayesian procedure that is able to incorporate prior knowledge of potential changes, but also computationally scalable. First, it requires only $3K$ registers for retaining relevant information

of $(R_{k,t}, L_{k,t}, \tilde{R}_{k,t})$ about the K local processes: the first two are on the observed data regarding the local change, and the last is on the prior knowledge of the local change. Second, since the two-dimensional local statistics $(R_{k,t}, L_{k,t})$ can be computed recursively and the “initial” values $\tilde{R}_{k,t}$ can be sampled directly from the prior distribution G at each time step, the computational cost of our TSSRP method is linear to the number K of local data streams. Thus our method can be easily implemented for real-time monitoring.

3.2 Choice of Parameters

The TSSRP algorithm involves several parameters. Below we will discuss the choice of these parameters.

Choice of the prior distribution: In practice, we could choose the prior distribution according to our prior knowledge. For example, in the manufacturing process, we may know that certain production lines could have a higher chance of being out of control. Alternatively, if no prior knowledge is present, we could choose some non-informative priors such as the uniform distribution or the point mass 0 distribution, i.e., $P(\tilde{R}_{k,t} = 0) = 1$. In the latter case, it reduces to a greedy sampling algorithm without randomization. In our numerical simulation studies in Section 4, we compare the performance of TSSRP with four different priors. The results suggest that the TSSRP significantly reduces the average detection delay regardless of the choice of the priors and that a valid prior can further improve the performance.

We note that Pollak (1985) also investigates the choice of the prior distribution G , but under a different context in which the randomized Shiryaev-Roberts statistics leads to an

equalizer stopping time T in that sense that $\mathbb{E}_\nu(T - \nu | T \geq \nu)$ is constant as a function of candidate change-point ν . Unfortunately, it is generally challenging to find an explicit solution for such prior distribution. Nevertheless, our purpose of the randomization is different from Pollak (1985): ours is for balancing the exploration and exploitation, while theirs is for almost minimax property.

Choice of r : Intuitively, the tuning parameter r in the stopping time (10) decides how many local sensors should be involved in the final decision making. Thus an ideal choice should be a plausible approximation of the actual number of changed data streams. If r is much smaller than the actual number of changed data streams, our final decision will not use all information that the data might provide. If r is much larger, our final decision will involve unnecessary noisy local statistics and lead to poor performance. Meanwhile, from the in-control performance viewpoint, as discussed in Wu (2019), the Shiryaev-Roberts statistic is heavy-tailed under the in-control hypothesis, and thus there is no practice difference between a smaller value of r (e.g., $r = 3$) to a larger value of r (e.g., $r = q$). Thus, when the total number of changed data streams is unknown, one might simply choose $r = q$.

Choice of A : The parameter A is the stopping threshold that controls the average run length to false alarm of our method, which is analogous to controlling the type I error. In practice, the threshold A is usually determined by Monte-Carlo simulation. One often uses the bisection method to find the smallest A so that the proposed method satisfies the false alarm constraint in (2).

3.3 Theoretical Properties of TSSRP

In this subsection, we provide some theoretical properties of the TSSRP algorithm: Theorem 1 establishes the close relationship between our algorithm and the Bayesian procedures. Theorem 2 investigates the in-control average run length properties, whereas Theorem 3 and Theorem 4 provide a deep understanding of the sensor layouts.

First, the following theorem provides theoretical bases for the proposed adaptive sampling in our TSSRP algorithm.

Theorem 1. *Assume that the change time ν_k for the k -th data stream has a prior $\text{Geometric}(p)$ distribution in (6), where the initial prior probability $\Pi_{k,0}$ has a prior G_p distribution. Suppose that $G_p/p \rightarrow G$ in distribution, then $R_{k,t}^*$ in (9) with its random component $\tilde{R}_{k,t} \sim G$ has the same distribution as*

$$\lim_{p \rightarrow 0} \frac{\Pi_{k,t}}{p(1 - \Pi_{k,t})}, \quad (11)$$

where $\Pi_{k,t} = \mathbf{P}(\nu_k \leq t | X_{k,1}^*, \dots, X_{k,t}^*)$ is the posterior estimation how likely a local change occurs, and $X_{k,t}^* = X_{k,t} \delta_{k,t}$ denotes the observed data.

The detailed proof of Theorem 1 is presented in the Appendix A.1 in online supplementary material. By this theorem, sampling based on the largest values of local statistics $R_{k,t}^*$ in (9) is mathematically equivalent to sampling based on the largest values of random realizations of the posterior distribution $\Pi_{k,t}$ as $p \rightarrow 0$, since $\frac{u}{p(1-u)}$ is a monotonic increasing function of u . Hence, the adaptive sampling policy in the TSSRP algorithm is the limit of the Thompson Sampling policy.

Second, we investigate the ARL to false alarm of the proposed TSSRP algorithm in the following theorem.

Theorem 2 (Average Run Length to False Alarm). $\mathbb{E}_\infty(T) \geq A/K$. Moreover, $\mathbb{E}_\infty(T) = O(A)$, where $O(A)/A$ is bounded as $A \rightarrow \infty$.

Theorem 2 provides us guidance to select conservative upper and lower bounds of A . Specifically, $K \cdot ARL$ can serve as the upper bound in the bisection search to speed up the threshold choosing procedure. The detailed proof of Theorem 2 is given in Appendix A.2 in online supplementary material. Unfortunately, it remains an open problem to derive the bounds on the average detection delays.

Next, we investigate the properties for sensor layouts. Theorem 3 shows that the sensor layout will go through all the data streams eventually when the system is in control.

Theorem 3. *Let S_t be the sensor layout at time t . Then under $H_0 : \nu = \infty$, there exists a $t' > t$ such that $P(k \in S_{t'}) > 0$, for each $t > 0$ and each $k, 1 \leq k \leq K$.*

Theorem 3 implies that each variable has a chance to be explored, regardless of the sensor deployments in previous steps when no changes occur. In other words, the algorithm performs similar to random sampling in the in-control state.

Finally, Theorem 4 below implies that the sensor layout of the TSSRP algorithm will eventually converge to the changed data streams when the system is out of control.

Theorem 4. *Let S_t be the sensor layout at time t . Then under $H_1 : \nu < \infty$, we have $P(k \in S_t, \forall t > t_0 | k \in S_{t_0}) > 0$ for all changed data stream k , and all $t_0 > \nu$.*

We have shown that the sensor layout will not stay on any unchanged data streams forever in Theorem 3. The sensors will eventually be redistributed to the affected data streams at a certain time. Theorem 4 states that once a sensor is deployed to the out-of-control data stream, then there is a nonzero probability that the sensor will stay on this

data stream forever. This property ensures that the sensors will eventually keep monitoring the affected data streams. The proofs of Theorem 3 and Theorem 4 are given in Appendix A.3 in online supplementary material.

In summary, our TSSRP algorithm admits a nice property that it does more exploration when the system is in control, while more exploitation when the system is out of control. Intuitively, such nice property comes from the structure of the local statistics. When we observe a new data, the recursive formula for local statistics involves the likelihood ratio $f_{\theta_{k,1}}(x_{k,t})/f_{\theta_{k,0}}(x_{k,t})$. Under the in-control state, the expectation of the increment of all local statistics at time t is 1, i.e., $\mathbb{E}[R_{k,t}^*] = t, \forall k, t$. Thus our adaptive sampling is similar to random sampling under the in-control state. Under the out-of-control state, the likelihood ratio under the post-change distribution will be more likely to be greater than one. Therefore, the affected data streams' local statistics will increase faster than those of unaffected data streams, which enables the sensor layout to converge to those affected local components.

4 Simulation Experiments

In this section, we report the numerical performance of the proposed TSSRP algorithm and compare it with the existing algorithms. The general setting for our simulations is as follows. We consider monitoring $K = 100$ independent data streams. We assume that $q = 10$ out of $K = 100$ data streams can be monitored at each time step. The nominal value for the ARL to a false alarm is fixed as $\mathbb{E}_{\infty}T = \gamma = 1000$. We follow the classic approach (Xie and Siegmund, 2013; Xie et al., 2013; Mei, 2010) and report the average

detection delay when the change occurs at time $\nu = 1$. It is the most challenging setting because it is more difficult to detect when the change occurred at the very beginning.

We report our simulation results in two subsections, depending on different true generative models of the data. Subsection 4.1 focuses on the statistical efficiency of the TSSRP algorithm when our prior knowledge on the candidate affected local streams and the Gaussian distribution of the data are valid, and Subsection 4.2 considers the robustness of our algorithm under the mis-specified models. In our simulation studies below, all numerical results are based on 1000 Monte Carlo replications.

For our proposed TSSRP algorithm, we consider four choices for the prior distribution G on the initial statistics $\tilde{R}_{k,t=0}$, which is the prior knowledge of how likely a local stream might be affected by the change:

- (i) G_0 : Uniform $U[0.5, 1]$ for the first ten local streams, and Uniform $U[0, 0.5]$ for the remaining local streams;
- (i) G_1 : Uniform $U[0.5, 1]$ for the first five local streams, and Uniform $U[0, 0.5]$ for the remaining local streams;
- (ii) G_2 : Uniform $U[0, 1]$ for all local streams;
- (iii) G_3 : the point mass 0 for all local streams, i.e., $P_0 = P(\tilde{R}_{k,t} = 0) = 1$.

We referred them to as TSSRP(G_0), TSSRP(G_1), TSSRP(G_2) and TSSRP(G_3), respectively.

We compare our TSSRP algorithms with the baseline Top-r Based Adaptive Sampling (TRAS) algorithm proposed by Liu et al. (2015). The TRAS algorithm first constructs a

local CUSUM statistic for each observed variable. For the unobserved variables, the local statistics are updated by adding a compensation parameter Δ . That is, under our notation, each local stream computes a local statistic

$$W_{k,t} = \begin{cases} \max(W_{k,t-1} + \log \frac{f_{\theta_{k,1}}(X_{k,t})}{f_{\theta_{k,0}}(X_{k,t})}, 0), & \text{if } \delta_{k,t} = 1; \\ w_{k,t-1} + \Delta, & \text{if } \delta_{k,t} = 0. \end{cases} \quad (12)$$

where Δ is the so-called compensation parameter for unobserved data streams. Next, the TRAS algorithm combines the top- r local statistics to determine whether to raise an alarm, i.e., the stopping time of the TRAS algorithm is given by

$$\tau(a) = \inf\{t \geq 1 : \sum_{k=1}^r W_{(k),t} \geq a\}, \quad (13)$$

where $W_{(k),t}$'s are the order statistics of $W_{k,t}$'s in (12). Moreover, the TRAS algorithm adaptively deploys the sensors to the data streams with q largest local statistics $W_{k,t}$'s in (12) at the next time step $t+1$. As reported in Liu et al. (2015), the average detection delay performance of the TRAS algorithm is sensitive to the choice of the compensation parameter Δ , and it remains an open problem to decide how to choose it suitably. In our simulations below, we present results obtained by three different choices of $\Delta = 0.03, 0.05, 0.1$.

4.1 Statistical Efficiency

In this subsection, we focus on the statistical efficiency of the TSSRP algorithms when the prior knowledge is valid. We consider the scenario where we monitor $K = 100$ independent Gaussian data streams whose pre-change distributions are $N(0, 1)$, and the first r local streams change to the post-change distribution $N(1.5, 1)$. We compare our TSSRP algorithm against other procedures with the four choices of initial distribution G and the

correct post-change mean $\mu_1 = 1.5$. We vary the number of changed data streams ranging from 1 to 10. As mentioned in Section 3.2, the parameter r in the global stopping time defined in (10) ideally should be the number of changed data streams. The latter, however, is usually unknown in practice. We report the simulation results of different global monitoring schemes under a large $r = 10$ in Table 1. Additional simulation results under a small $r = 3$ is deferred to Table 5 in supplementary material. The corresponding standard errors are also included in these tables to characterize the average detection delay distribution.

Table 1: average detection delay under various number of changed data streams for the evaluation of the statistical efficiency experiments when the data is independent multivariate Gaussian distributed. All the experiments are conducted under $r = 10$.

The number of changes	1	3	5	8	10
TSSRP(G_0)	12.15(0.23)	7.67(0.07)	6.66(0.05)	6.05(0.04)	5.81(0.03)
TSSRP(G_1)	12.06(0.23)	7.59(0.07)	6.75(0.05)	6.57(0.04)	6.49(0.04)
TSSRP(G_2)	18.84(0.33)	11.93(0.14)	10.05(0.11)	8.67(0.08)	8.22(0.07)
TSSRP(G_3)	19.43(0.35)	11.79(0.14)	9.84(0.11)	8.74(0.08)	8.04(0.07)
TRAS($\Delta = 0.03$)	36.12(0.60)	21.10(0.25)	17.01(0.20)	13.43(0.15)	11.87(0.13)
TRAS($\Delta = 0.05$)	36.79(0.54)	22.84(0.24)	18.52(0.18)	15.17(0.13)	13.52(0.12)
TRAS($\Delta = 0.1$)	63.43(0.44)	37.87(0.25)	30.47(0.18)	25.39(0.13)	22.89(0.12)
RSADA	71.87(1.63)	36.61(0.69)	26.94(0.51)	21.20(0.38)	18.54(0.34)

We make two key observations from Table 1. First, all four variants of our proposed TSSRP algorithms are statistically efficient in the sense of having significantly smaller av-

erage detection delays compared to other procedures. When incorporating the correct prior information, i.e., with the prior G_0 and G_1 on the initial statistics $\tilde{R}_{k,t=0}$, the TSSRP algorithm achieves the smallest average detection delays, since it appropriately incorporates the Bayesian information on the spatial locations of changes. Even if we use non-informative priors such as G_2 or G_3 , the TSSRP algorithm still provides good performance as compared to the baseline TRAS algorithm, suggesting that the choice of priors in the TSSRP algorithm can be flexible. Second, fewer tuning efforts are required for the TSSRP algorithm, because we set the likelihood of the unobserved or missing data to be 1 under the Bayesian framework. Moreover, the performance of the TSSRP algorithm is relatively stable to the tuning parameter r , e.g., the number of local sensors involved in the final decision making, see table 5 in the supplementary material. This is consistent with our intuition that the Shiryaev-Roberts statistics is the exponent of the CUSUM statistics, and the sum of top- r Shiryaev-Roberts statistics mainly captures the maximum local statistics across all the data streams, since $\sum_{i=1}^r \exp(a_i) \sim \exp(a_1)$ for large values of ordered sequences $a_1 > a_2 > \dots > a_r$. Therefore, our TSSRP algorithm is not only statistically efficient but also easy to tune and use in practice.

4.2 Robustness

In this subsection, we focus on the robustness of the TSSRP algorithm when the underlying generative model is mis-specified. We focus on two cases: (i) TSSRP algorithm is constructed for the post-change mean lower bound $\mu_1 = 1.5$ when the true post-change mean of Gaussian distribution is $\mu_{1,true} = 2$ and (ii) TSSRP algorithm is constructed for

Table 2: average detection delay under various number of changed data streams for robustness experiments when the post-change parameter is mis-specified. The data streams follows Gaussian distribution.

The number of changes	1	3	5	8	10
TSSRP(G_0)	7.37(0.10)	5.43(0.03)	4.98(0.03)	4.54(0.02)	4.43(0.02)
TSSRP(G_1)	7.33(0.10)	5.33(0.03)	4.87(0.03)	4.77(0.03)	4.72(0.02)
TSSRP(G_2)	8.64(0.17)	5.84(0.07)	5.64(0.06)	5.49(0.05)	5.32(0.04)
TSSRP(G_3)	12.77(0.18)	8.28(0.09)	7.18(0.07)	6.16(0.05)	5.87(0.05)
TRAS($\Delta = 0.03$)	27.03(0.42)	16.42(0.21)	12.69(0.15)	10.03(0.11)	8.87(0.10)
TRAS($\Delta = 0.05$)	27.79(0.34)	17.42(0.18)	14.38(0.15)	11.10(0.11)	9.91(0.09)
TRAS($\Delta = 0.1$)	44.93(0.28)	27.73(0.17)	22.40(0.13)	18.78(0.11)	17.11(0.10)

Gaussian distributions when the data follows t distributions with degree of freedom $df = 5$.

The results for the first case are shown in Table 2, which summarizes the average average detection delay and the corresponding standard errors for TSSRP under four choices of priors and TRAS under $\Delta = 0.03, 0.05, 0.1$. As in the case in subsection 4.1, we see that TSSRP outperforms TRAS and that performance of TSSRP generally improves as the prior information gets correct. Moreover, compared with the results in Table 1, we see that a larger magnitude of change is easier to detect, suggesting that we can use the smallest magnitude of change to specify the post-change parameters if they are unknown.

For the latter case, we determine the stopping threshold A by 1000 Monte Carlo simulations under the t distribution with mean 0, and we choose $r = 10$ to construct the global

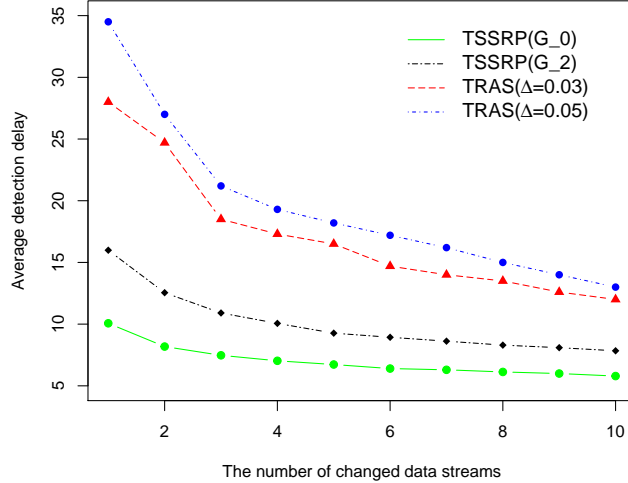


Figure 1: average detection delay versus the number of changed data streams for robustness experiments when the distribution is mis-specified. The data streams follows the t distribution.

statistics. Figure 1 plots the average detection delays against the number of changed data streams for TSSRP and TRAS, with varying priors (G_1 and G_2) and tuning parameters ($\Delta = 0.03, 0.05$). We observe a similar pattern as before, again reinforcing the usefulness of incorporating Bayesian information.

Our results suggest that our TSSRP algorithm can be robust when the distributional assumption is somewhat violated, and our TSSRP still outperforms the TRAS algorithm. This experiment further bolsters the stability of our algorithm, indicating that the hypothesized distributions can be slightly different from the underlying distributions of the data, and our algorithm will still raise an alarm quickly. In theory, the robustness of the TSSRP algorithm depends heavily on the robustness of the likelihood $\frac{f_1}{f_0}$. As a result, when the underlying models are significantly mis-specified, our TSSRP algorithm can be less efficient




Figure 2: Left: 2-D illustration of the hot forming process. Right: Bayesian network for the hot forming process

as compared to other robust methods such as the RSADA algorithm in Xian et al. (2021) that is based on ranks. In such a case, our algorithm can be extended to a more robust variant by considering the likelihood of sequential ranks as in Gordon and Pollak (1995), which is beyond the scope of this article, and we will leave it as future work.

5 Case Study

In this section, we evaluate the performance of the TSSRP algorithm on a real case example: the hot forming process. We also give an additional solar flare detection example on the high-dimensional case in Appendix B in supplementary material.

We consider the Hot Forming Process example in Li and Jin (2010). We want to detect anomalies in this physical system. Figure 2 (Left) illustrates a two-dimensional (2-D) physical illustration of the hot forming process. Li and Jin (2010) identified the causal relationship of the five variables in this process: the final dimension of workpiece X_1 , the

tension in workpiece X_2 , the material flow stress X_3 , temperature X_4 and Blank Holding Force X_5 , which can be represented as a Bayesian network. All the variables are proven to follow standard normal distribution when the system is under normal operating conditions. Li and Jin (2010) gave the parameterization model of the Bayesian network, and Figure 2 Right illustrates the dependence across the five variables:

$$X_i = \sum_{X_j \in p(X_i)} w_{j,i} X_j + \epsilon_i \quad (14)$$

where $p(X) = \{Y : Y \rightarrow X \in \text{Edge Set}\}$ denotes the parents of X ; and $w(X_j, X_i)$ is the weight of the edge $X_j \rightarrow X_i$, which refers to the causal influence from X_j to X_i ; $\epsilon_i \sim N(0, \sigma_i^2)$ is the independent Gaussian noise.

In this study, we assume that the Bayesian network is unknown to us. We will only use the network structure to generate data under different scenarios. We generate the changed root variables by setting the true mean change as two, and the remaining variables according to the Bayesian network. In the algorithm, we set the post-change mean as the interested smallest shift magnitude $\mu_{k,1} = 1.5$ according to the characteristics of the actual system. We set the in-control ARL to be 100 and $r = 2$ as in Liu et al. (2015). In each replicate, the changed data streams and the initial sensor layouts are selected randomly. We evaluate the average detection delay $D(T)$ as the average average detection delay under any change possibilities.

Table 3 summarizes the average detection delay comparisons between the TSSRP algorithm and the TRAS algorithm in the single change and two changes cases. It implies that the performance of our TSSRP algorithm is better than that of the baseline TRAS algorithm in this hot forming procedure.

Table 3: Comparisons of the average detection delay between the TSSRP algorithm and the TRAS algorithm under different numbers of changed root variables in the hot forming process example

The number of changes	1	2
TSSRP($U[0, 1]$)	5.84(0.10)	4.54(0.06)
TSSRP(P_0)	6.48(0.09)	5.29(0.07)
TRAS($\Delta = 0.01$)	8.41(0.13)	7.39(0.11)
TRAS($\Delta = 0.1$)	8.17(0.13)	7.09(0.12)
TRAS($\Delta = 0.5$)	9.94(0.15)	8.78(0.14)

We are interested in studying how the sensor layouts update over the time under the in-control and out-of-control states empirically to validate Theorem 3 and Theorem 4. We summarize the average percentages of each data stream being observed under the two states in Table 4, based on 1000 replicates. Specifically, it is defined as

$$\text{percentage of being observed} = \frac{\#\text{time steps being observed}}{\#\text{total time steps until raising an alarm}}.$$

Here, the out-of-control state is when X_1 and X_2 change at the very beginning. The simulations further confirm that our TSSRP algorithm works similar to random sampling when the system is in control, and greedily selects the changed data streams when the system is out-of-control.

Table 4: Sensor layouts distribution under the in-control and out-of-control states in the hot forming process example

Variable	Percentage (In-control)	Percentage (Out-of-control)
X_1	0.408	0.718
X_2	0.396	0.606
X_3	0.400	0.226
X_4	0.400	0.224
X_5	0.396	0.226

6 Conclusions and Discussion

Processing high-velocity streams of high-dimensional data in resource-constrained environments is a big challenge. In this paper, we propose a bandit change-point detection approach to adaptively sample useful local components and determine a global stopping time for real-time monitoring of high-dimensional streaming data. Our proposed algorithm, termed Thompson-Sampling Shiryaev-Roberts-Pollak (TSSRP) algorithm, can balance between exploiting those observed local components that maximize the immediate detection performance and exploring not-been-monitored local components that might accumulate new information to improve future detection performance. Our numerical simulations and case studies show that the TSSRP algorithm can significantly reduce the average detection delay compared to the existing methods.

This work can be extended in several directions. First, based on the numerical simulation studies, we conjecture that our proposed TSSRP algorithm is first-order asymptotically

optimal under a general setting. Still, it remains an open problem to prove it, as it is highly non-trivial to analyze the expected average detection delay of the proposed method. Second, instead of a fixed number of active sensors, one could consider changing the number of active sensors per time step, and increase the number of sensors if a change likely occurs. Third, it is also interesting to find an optimal value of the number of active sensors that can adaptively adjust to make the best use of the resource.

Finally, we remark that our TSSRP algorithm is a computationally scalable representation of the limit of Bayesian procedures under the simplest model assumption where the data is i.i.d. and the post-change parameters are known. It can be extended to handle the case when the post-change parameters are unknown if we introduce a prior distribution on the post-change parameters. Furthermore, for more complicated models where the data streams have a spatial or temporal correlation structure, the proposed Bayesian and Thompson sampling framework can still be applicable if we can update the posterior distribution efficiently, say, via Markov chain Monte Carlo (MCMC). Therefore, our work opens a new research direction on statistical process control and sequential change-point detection when monitoring high-dimensional data streams under the sampling control.

Supplementary Materials

In the supplementary materials, we provide (A) the detailed proofs of all theorems, (B) an additional case study in Solar Flare data, (C) additional simulation studies on our proposed algorithm when raising a global alarm based on the sum of the largest $r = 3$ local statistics, and (D) the comparison with more global decision policies. The zip file contains

R codes for our algorithm.

Acknowledgments

The authors thank the editor, the associate editor, and two reviewers for their invaluable comments that greatly help to improve the article.

Funding

W.Z. is supported in part by an ARC-TRIAD fellowship from the Georgia Institute of Technology, and a Computing Innovation Fellowship from the Computing Research Association (CRA) and the Computing Community Consortium (CCC). This work was completed while W.Z. was at Georgia Institute of Technology. Y.M. is supported in part by NSF grant DMS-2015405. W.Z. and Y.M. are also supported in part by the National Center for Advancing Translational Sciences of the National Institutes of Health under Award Number UL1TR002378. The content is solely the responsibility of the authors and does not necessarily represent the official views of the National Institutes of Health.

References

Agrawal, S. and Goyal, N. (2012). Analysis of thompson sampling for the multi-armed bandit problem. In *Proceedings of the 25th Annual Conference on Learning Theory*, pages 39.1–39.26.

- Agrawal, S. and Goyal, N. (2013). Thompson sampling for contextual bandits with linear payoffs. In *Proceedings of the 30th International Conference on Machine Learning*, pages 127–135.
- Anantharam, V., Varaiya, P., and Walrand, J. (1987). Asymptotically efficient allocation rules for the multiarmed bandit problem with multiple plays-Part I: IID rewards. *IEEE Transactions on Automatic Control*, 32(11):968–976.
- Bass, T. (1999). Multisensor data fusion for next generation distributed intrusion detection systems. In *Proceedings of the IRIS National Symposium on Sensor and Data Fusion*, volume 24, pages 24–27.
- Basseville, M. and Nikiforov, I. V. (1993). *Detection of Abrupt Changes: Theory and Application*. Prentice Hall Englewood Cliffs.
- Bubeck, S. and Cesa-Bianchi, N. (2012). Regret analysis of stochastic and nonstochastic multi-armed bandit problems. *Foundations and Trends® in Machine Learning*, 5(1):1–122.
- Cao, Y., Wen, Z., Kveton, B., and Xie, Y. (2019). Nearly optimal adaptive procedure with change detection for piecewise-stationary bandit. In *The 22nd International Conference on Artificial Intelligence and Statistics*, pages 418–427. PMLR.
- Chan, H. P. (2017). Optimal sequential detection in multi-stream data. *The Annals of Statistics*, 45(6):2736–2763.

- Chapelle, O. and Li, L. (2011). An empirical evaluation of thompson sampling. In *Advances in neural information processing systems*, pages 2249–2257.
- Cho, H. and Fryzlewicz, P. (2015). Multiple-change-point detection for high dimensional time series via sparsified binary segmentation. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 77(2):475–507.
- Chu, L. and Chen, H. (2019). Asymptotic distribution-free change-point detection for multivariate and non-euclidean data. *The Annals of Statistics*, 47(1):382–414.
- Ding, Y., Elsayed, E. A., Kumara, S., Lu, J.-C., Niu, F., and Shi, J. (2006). Distributed sensing for quality and productivity improvements. *IEEE Transactions on Automation Science and Engineering*, 3(4):344–359.
- Ghatak, G. (2020). A change-detection based Thompson Sampling framework for non-stationary bandits. *IEEE Transactions on Computers*, 70(10):1670–1676.
- Gittins, J., Glazebrook, K., and Weber, R. (2011). *Multi-Armed Bandit Allocation Indices*. John Wiley & Sons.
- Gordon, L. and Pollak, M. (1995). A robust surveillance scheme for stochastically ordered alternatives. *The Annals of Statistics*, 23(4):1350–1375.
- Kaufmann, E., Cappé, O., and Garivier, A. (2016). On the complexity of best-arm identification in multi-armed bandit models. *The Journal of Machine Learning Research*, 17(1):1–42.

- Lai, T. L. (1995). Sequential changepoint detection in quality control and dynamical systems. *Journal of the Royal Statistical Society: Series B (Methodological)*, 57(4):613–644.
- Lai, T. L. (1998). Information bounds and quick detection of parameter changes in stochastic systems. *IEEE Transactions on Information Theory*, 44(7):2917–2929.
- Lai, T. L. and Robbins, H. (1985). Asymptotically efficient adaptive allocation rules. *Advances in Applied Mathematics*, 6(1):4–22.
- Li, J. (2019). A two-stage online monitoring procedure for high-dimensional data streams. *Journal of Quality Technology*, 51(4):392–406.
- Li, J. (2020). Efficient global monitoring statistics for high-dimensional data. *Quality and Reliability Engineering International*, 36(1):18–32.
- Li, J. and Jin, J. (2010). Optimal sensor allocation by integrating causal models and set-covering algorithms. *IIE Transactions*, 42(8):564–576.
- Li, W., Xiang, D., Tsung, F., and Pu, X. (2020). A diagnostic procedure for high-dimensional data streams via missed discovery rate control. *Technometrics*, 62(1):84–100.
- Liu, K., Mei, Y., and Shi, J. (2015). An adaptive sampling strategy for online high-dimensional process monitoring. *Technometrics*, 57(3):305–319.
- Liu, K., Zhang, R., and Mei, Y. (2019). Scalable sum-shrinkage schemes for distributed monitoring large-scale data streams. *Statistica Sinica*, 29:1–22.

- Lorden, G. (1971). Procedures for reacting to a change in distribution. *The Annals of Mathematical Statistics*, 42(6):1897–1908.
- Mei, Y. (2010). Efficient scalable schemes for monitoring a large number of data streams. *Biometrika*, 97(2):419–433.
- Mei, Y. (2011). Quickest detection in censoring sensor networks. In *Information Theory Proceedings (ISIT), 2011 IEEE International Symposium on*, pages 2148–2152. IEEE.
- Nabhan, M., Mei, Y., and Shi, J. (2021). Correlation-based dynamic sampling for online high-dimensional process monitoring. *Journal of Quality Technology*, 53(3):289–308.
- Pandelis, D. G. and Teneketzis, D. (1999). On the optimality of the gittins index rule for multi-armed bandits with multiple plays. *Mathematical Methods of Operations Research*, 50(3):449–461.
- Pollak, M. (1985). Optimal detection of a change in distribution. *The Annals of Statistics*, 13(1):206–227.
- Pollak, M. (1987). Average run lengths of an optimal method of detecting a change in distribution. *The Annals of Statistics*, 15(2):749–779.
- Poor, H. V. and Hadjiliadis, O. (2008). *Quickest Detection*. Cambridge University Press.
- Robbins, H. (1952). Some aspects of the sequential design of experiments. *Bulletin of the American Mathematical Society*, 58(5):527–535.
- Robbins, H. (1985). Some aspects of the sequential design of experiments. In *Herbert Robbins Selected Papers*, pages 169–177. Springer.

- Roberts, S. (1966). A comparison of some control chart procedures. *Technometrics*, 8(3):411–430.
- Scott, S. L. (2010). A modern bayesian look at the multi-armed bandit. *Applied Stochastic Models in Business and Industry*, 26(6):639–658.
- Shiryaev, A. N. (1963). On optimum methods in quickest detection problems. *Theory of Probability & Its Applications*, 8(1):22–46.
- Tartakovsky, A., Nikiforov, I., and Basseville, M. (2014). *Sequential Analysis: Hypothesis Testing and Changepoint Detection*. Chapman and Hall/CRC.
- Thompson, W. R. (1933). On the likelihood that one unknown probability exceeds another in view of the evidence of two samples. *Biometrika*, 25(3/4):285–294.
- Viswanath, B., Bashir, M. A., Crovella, M., Guha, S., Gummadi, K. P., Krishnamurthy, B., and Mislove, A. (2014). Towards detecting anomalous user behavior in online social networks. In *23rd USENIX security symposium (USENIX security 14)*, pages 223–238.
- Wang, A., Xian, X., Tsung, F., and Liu, K. (2018). A spatial-adaptive sampling procedure for online monitoring of big data streams. *Journal of Quality Technology*, 50(4):329–343.
- Wang, Y. and Mei, Y. (2015). Large-scale multi-stream quickest change detection via shrinkage post-change estimation. *IEEE Transactions on Information Theory*, 61(12):6926–6938.
- Wu, Y. (2019). A combined SR-CUSUM procedure for detecting common changes in panel data. *Communication in Statistics: Theory and Methods*, 48(17):4302–4319.

- Xian, X., Wang, A., and Liu, K. (2018). A nonparametric adaptive sampling strategy for online monitoring of big data streams. *Technometrics*, 60(1):14–25.
- Xian, X., Zhang, C., Bonk, S., and Liu, K. (2021). Online monitoring of big data streams: A rank-based sampling algorithm by data augmentation. *Journal of Quality Technology*, 53(2):135–153.
- Xie, Y., Huang, J., and Willett, R. (2013). Change-point detection for high-dimensional time series with missing data. *IEEE Journal of Selected Topics in Signal Processing*, 7(1):12–27.
- Xie, Y. and Siegmund, D. (2013). Sequential multi-sensor change-point detection. *The Annals of Statistics*, 41(2):670–692.
- Yang, S., Santillana, M., and Kou, S. C. (2015). Accurate estimation of influenza epidemics using google search data via argo. *Proceedings of the National Academy of Sciences*, 112(47):14473–14478.
- Zhang, C. and Hoi, S. C. H. (2019). Paritally observable multi-sensor sequential change detection: A combinatorial multi-armed bandit approach. In *Proceedings of the Thirty-Third AAAI Conference on Artificial Intelligence (AAAI-19)*, volume 33, pages 5733–5740.
- Zhang, N. R. and Siegmund, D. O. (2012). Model selection for high-dimensional, multi-sequence change-point problems. *Statistica Sinica*, 22(4):1507–1538.
- Zhao, Q. (2019). *Multi-Armed Bandits: Theory and Applications to Online Learning in*

Networks (Synthesis Lectures on Communication Networks). Morgan & Claypool Publishers.

Zou, C. and Qiu, P. (2009). Multivariate statistical process control using lasso. *Journal of the American Statistical Association*, 104(488):1586–1596.

Zou, C., Wang, Z., Zi, X., and Jiang, W. (2015). An efficient online monitoring method for high-dimensional data streams. *Technometrics*, 57(3):374–387.