

RESEARCH

Open Access

Sketching for sequential change-point detection



Yang Cao¹, Andrew Thompson², Meng Wang³ and Yao Xie^{1*}

Abstract

We present sequential change-point detection procedures based on linear sketches of high-dimensional signal vectors using generalized likelihood ratio (GLR) statistics. The GLR statistics allow for an unknown post-change mean that represents an anomaly or novelty. We consider both fixed and time-varying projections, derive theoretical approximations to two fundamental performance metrics: the average run length (ARL) and the expected detection delay (EDD); these approximations are shown to be highly accurate by numerical simulations. We further characterize the relative performance measure of the sketching procedure compared to that without sketching and show that there can be little performance loss when the signal strength is sufficiently large, and enough number of sketches are used. Finally, we demonstrate the good performance of sketching procedures using simulation and real-data examples on solar flare detection and failure detection in power networks.

Keywords: Change-point detection, Streaming data, Sketching, Anomaly detection, Statistical signal processing

1 Introduction

Online change-point detection from high-dimensional streaming data is a fundamental problem arising from applications such as real-time monitoring of sensor networks, computer network anomaly detection, and computer vision (e.g., [2, 3]). To reduce data dimensionality, a conventional approach is *sketching* (see, e.g., [4]), which performs random projection of the high-dimensional data vectors into lower-dimensional ones. Sketching has now been widely used in signal processing and machine learning to reduce dimensionality and algorithm complexity and achieve various practical benefits [5–11].

We consider change-point detection using linear sketches of high-dimensional data vectors. Sketching reduces the computational complexity of the detection statistic from $\mathcal{O}(N)$ to $\mathcal{O}(M)$, where N is the original dimensionality and M is the dimensionality of sketches. Since we would like to perform real-time detection, any reduction in computational complexity (without incurring much performance loss) is highly desirable. Sketching

also offers practical benefits. For instance, for large sensor networks, it reduces the burden of data collection and transmission. It may be impossible to collect data from all sensors and transmit them to a central hub in real time, but this can be done if we only select a small subset of sensors to collect data at each time. Sketching also reduces data storage requirement. For instance, change-point detection using the generalized likelihood ratio statistic, although robust, is non-recursive. Thus, one has to store historical data. Using sketching, we only need to store the much lower dimensional sketches rather than the original high-dimensional vectors.

In this paper, we present a new sequential *sketching procedure* based on the generalized likelihood ratio (GLR) statistics. In particular, suppose we may choose an $M \times N$ matrix A with $M \ll N$ to linearly project the original data: $y_t = Ax_t, t = 1, 2, \dots$. Assume the pre-change vector is zero-mean Gaussian distributed and the post-change vector is Gaussian distributed with an *unknown* mean vector μ while the covariance matrix is unchanged. Here, we assume the mean vector is unknown since it typically represents an anomaly. The GLR statistic is formed by replacing the unknown μ with its maximum likelihood ratio estimator (e.g., [12]). Then we further generalize to the setting with time-varying projections A_t of dimension $M_t \times N$. We demonstrate the good performance of our

*Correspondence: yao.xie@isye.gatech.edu

This paper was presented [in part] at the GlobalSIP 2015 [1].

¹H. Milton Stewart School of Industrial and Systems Engineering, Georgia Institute of Technology, Atlanta, GA, USA

Full list of author information is available at the end of the article

procedures by simulations, a real data example of solar flare detection, and a synthetic example of power network failure detection with data generated using real-world power network topology.

1.1 Our contribution

Our theoretical contribution is mainly in two aspects. We obtain analytic expressions for two fundamental performance metrics for the sketching procedures: the average run length (ARL) when there is no change and the expected detection delay (EDD) when there is a change-point, for both fixed and time-varying projections. Our approximations are shown to be highly accurate using simulations. These approximations are quite useful in determining the threshold of the detection procedure to control false alarms, without having to resort to the onerous numerical simulations. Moreover, we characterize the relative performance of the sketching procedure compared to that without sketching. We examine the EDD ratio when the sketching matrix A is either a random Gaussian matrix or a sparse 0-1 matrix (in particular, an expander graph). We find that, as also verified numerically, when the signal strength and M are sufficiently large, the sketching procedure may have little performance loss. When the signal is weak, the performance loss can be large when M is too small. In this case, our results can be used to find the minimum M such that performance loss is bounded, assuming certain worst case signal and for a given target ARL value.

To the best of our knowledge, our work is the first to consider sequential change-point detection using the generalized likelihood ratio statistic, assuming the *post-change mean is unknown to represent an anomaly*. The only other work [13] that considers change-point detection using linear projections assumes the post-change mean is known and further to be sparse. Our results are more general since we do not make such assumptions. Assuming the post-change mean to be unknown provides a more useful procedure since in change-point detection, the post-change setup is usually unknown. Moreover, [13] considers Shiryaev-Robert's procedure, which is based on a different kind of detection statistic than the generalized likelihood ratio statistic considered here. The theoretical analyses therein consider slightly different performance measures, the probability of false alarm, and average detection delay, and our analyses are completely different.

Our work is also distinctive from the existing Statistical Process Control (SPC) charts using random projections (reviewed below in Section 1.3) in that (1) we developed new theoretical results for the sequential GLR statistic, (2) we consider the sparse 0-1 and time-varying projections, and (3) we study the amount of dimensionality reduction

can be performed (i.e., the minimum M) such that there is little performance loss.

1.2 Notations and outline

Our notations are standard: χ_k^2 denotes the chi-square distribution with degree-of-freedom k ; I_n denotes an identity matrix of size n ; X^\dagger denotes the pseudo-inverse of a matrix X ; $[\mathbf{x}]_i$ denotes the i th coordinate of a vector \mathbf{x} ; $[X]_{ij}$ denotes the ij th element of a matrix X ; and \mathbf{x}^\top denotes the transpose of a vector or matrix \mathbf{x} .

The rest of the sections are organized as follows. We first review some related work. Section 2 sets up the formulation of the sketching problem for sequential change-point detection. Section 3 presents the sketching procedure. Section 4 contains the performance analysis of the sketching procedures. Section 5 and Section 6 demonstrate good performance of our sketching procedures using simulation and real-world examples. Section 7 concludes the paper. All proofs are delegated to the appendix.

1.3 Related work

In this paper, we use the term “sketching” in a broader sense to mean that our observations are linear projections of the original signals. We are concerned with how to perform sequential change-point detection using these linear projections. The traditional sketching [4] is concerned with designing linear dimensionality reduction techniques to solve the inverse linear problem $A\mathbf{x} = \mathbf{b}$, where \mathbf{b} is of greater dimension than \mathbf{x} . This can be cast as a problem of designing a dimensionality reduction (sketching) matrix S such that $S\mathbf{b} = SA\mathbf{x}$ is of smaller dimension to reduce computational cost. In our problem, the linear projections can be designed or they can be determined by problem set-up (such as missing data or subsampling procedure).

A closely related work is [14], which considers one-dimensional observations, and the pre-change distribution is Gaussian with zero mean and unit variance, and the post-change distribution is Gaussian with *unknown* mean and unit variance. Siegmund and Venkatraman [14] also uses the GLR statistic, by estimating the post-change mean and plug-in the likelihood ratio statistic. Our strategy for deriving the detection statistic is similar to [14]. However, there is one crucial difference. Since the number of linear projections is much smaller than the original dimension, we cannot obtain a unique MLE for the post-change mean vector, but can only determine the equation that the MLE needs to satisfy. Thus, the derivation and the analysis of GLR detection statistic for our setting are different from [14]. Another closely related work is [15]: we adapt a result therein (Theorem 1) in deriving the ARL and EDD of the sketching procedure, when the projection matrix is fixed. The scope of the two papers are quite different: [15] studies change-point detection when the post-change mean is sparse, and here, we are not

concerned with detecting sparse change but with detecting using linear projections of the original data; moreover, our analysis for the time-varying projection case is new and different from [15].

Change-point detection problems are closely related to industrial quality control and multivariate statistical process control (SPC) charts, where an observed process is assumed initially to be in control and at a change-point becomes out of control. The idea of using random projections for change detection has been explored for SPC in the pioneering work [16] based on U^2 multivariate control chart, the follow-up work [17] for cumulative sum (CUSUM) control chart and the exponential weighting moving average (EWMA) schemes, and in [18, 19] based on the Hotelling statistic. These works provide a complementary perspective from SPC design, while our method takes a different approach and is based on sequential hypothesis testing, treating both the change-point location and the post-change mean vector as unknown parameters. By treating the change-point location as an unknown parameter when deriving the detection statistic, the sequential hypothesis testing approach overcomes the drawback of some SPC methods due to a lack of memory, such as the Shewhart chart and the Hotelling chart, since they cannot utilize the information embedded in the entire sequence [20]. Moreover, our sequential GLR statistic may be preferred over the CUSUM procedure in the setting when it is difficult to specify the post-change mean vector.

This paper extends on our preliminary work reported in [1] with several important extensions. We have added (1) time-varying sketching projections and their theoretical analysis, (2) extensive numerical examples to verify our theoretical results, and (3) new real-data examples of solar flare detection and power failure detection.

Our work is related to compressive signal processing [21], where the problem of interest is to estimate or detect (in the fixed-sample setting) a sparse signal using compressive measurements. In [22], an offline test for a non-zero vector buried in Gaussian noise using linear measurements is studied; interestingly, a conclusion similar to ours is drawn that the task of detection within this setting is much easier than the tasks of estimation and support recovery. Another related work is [23], which considers a problem of identifying a subset of data streams within a larger set, where the data streams in the subset follow a distribution (representing anomaly) that is different from the original distribution; the problem considered therein is not a sequential change-point detection problem as the “change-point” happens at the onset ($t = 1$). In [24], an offline setting is considered and the goal is to identify k out of n samples whose distributions are different from the normal distribution f_0 . They use a “temporal” mixing of the samples over the finite time horizon. This

is different from our setting since we project over the signal dimension at each time. Other related works include kernel methods [25, 26] that focus on offline change-point detection. Finally, detecting transient changes in power systems has been studied in [27].

Another common approach to dimensionality reduction is principal component analysis (PCA) [28], which achieves dimensionality reduction by projecting the signal along the singular space of the leading singular values. In this case, \mathbf{A} or \mathbf{A}_t corresponds to the signal singular space. Our theoretical approximation for ARL and EDD can also be applied in these settings. It may not be easy to find the signal singular space when the dimensionality is high, since computing singular value decomposition can be expensive [29].

2 Formulation

2.1 Change-point detection as sequential hypothesis test

Consider a sequence of observations with an open time horizon $\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_t$, $t = 1, 2, \dots$, where $\mathbf{x}_t \in \mathbb{R}^N$ and N is the signal dimension. Initially, the observations are due to noise. There can be a time κ such that an unknown change-point occurs and it changes the mean of the signal vector. Such a problem can be formulated as the following hypothesis test:

$$\begin{aligned} H_0 : \mathbf{x}_t &\sim \mathcal{N}(0, \mathbf{I}_N), & t = 1, 2, \dots \\ H_1 : \mathbf{x}_t &\sim \mathcal{N}(0, \mathbf{I}_N), & t = 1, 2, \dots, \kappa, \\ &\mathbf{x}_t \sim \mathcal{N}(\boldsymbol{\mu}, \mathbf{I}_N), & t = \kappa + 1, \kappa + 2, \dots \end{aligned} \quad (1)$$

where the *unknown* mean vector is defined as

$$\boldsymbol{\mu} \triangleq [\mu_1, \dots, \mu_N]^T \in \mathbb{R}^N.$$

Without loss of generality, we have assumed the noise variance is 1. Our goal is to detect the change-point as soon as possible after it occurs, subjecting to the false alarm constraint. Here, we assume the covariance of the data to be an identity matrix and the change only happens to the mean.

To reduce data dimensionality, we linearly project each observation \mathbf{x}_t into a lower dimensional space, which we refer to as *sketching*. We aim to develop procedures that can detect a change-point using the low-dimensional *sketches*. In the following, we consider two types of linear sketching: the fixed projection and the time-varying projection.

Note that when the covariance matrix is known, the general problem is equivalent to (1), due to the following simple argument. Suppose we have the following hypothesis test:

$$\begin{aligned} H_0 : \mathbf{x}'_t &\sim \mathcal{N}(0, \boldsymbol{\Sigma}), & t = 1, 2, \dots \\ H_1 : \mathbf{x}'_t &\sim \mathcal{N}(0, \boldsymbol{\Sigma}), & t = 1, 2, \dots, \kappa, \\ &\mathbf{x}'_t \sim \mathcal{N}(\boldsymbol{\mu}', \boldsymbol{\Sigma}), & t = \kappa + 1, \kappa + 2, \dots \end{aligned} \quad (2)$$

where the covariance matrix Σ is positive definite. Denote the eigen-decomposition as $\Sigma = Q\Lambda Q^T$. Now, transform each observation x_t using $x_t = \Lambda^{-1/2}Q^T x'_t$, $t = 1, 2, \dots$, where $\Lambda^{-1/2}$ is a diagonal matrix with the diagonal entries being the inverse of the square root of the diagonal entries of Λ . Then, the original hypothesis test can be written in the same form as (1), by defining $\mu = \Lambda^{-1/2}Q^T \mu'$.

Fixed projection. Choose an $M \times N$ (possibly random) projection matrix \mathbf{A} with $M \ll N$. We obtain low-dimensional sketches via:

$$\mathbf{y}_t \triangleq \mathbf{A}\mathbf{x}_t, \quad t = 1, 2, \dots \quad (3)$$

Then the hypothesis test for the original problem (1), becomes the following hypothesis test based on the sketches (3)

$$\begin{aligned} H_0 : \mathbf{y}_t | \mathbf{A} &\sim \mathcal{N}(0, \mathbf{A}\mathbf{A}^T), \quad t = 1, 2, \dots \\ H_1 : \mathbf{y}_t | \mathbf{A} &\sim \mathcal{N}(0, \mathbf{A}\mathbf{A}^T), \quad t = 1, 2, \dots, \kappa, \\ &\mathbf{y}_t | \mathbf{A} \sim \mathcal{N}(\mathbf{A}\boldsymbol{\mu}, \mathbf{A}\mathbf{A}^T), \quad t = \kappa + 1, \kappa + 2, \dots \end{aligned} \quad (4)$$

Above, the distributions for the sketches are for given projections. Note that both mean and covariance structures are affected by the projections \mathbf{A} .

Time-varying projection. In certain applications, one may use different sketching matrices at each time. The projections are denoted by $\mathbf{A}_t \in \mathbb{R}^{M_t \times N}$ and the number of sketches M_t can change as well. The hypothesis test for sketches becomes:

$$\begin{aligned} H_0 : \mathbf{y}_t | \mathbf{A}_t &\sim \mathcal{N}(0, \mathbf{A}_t \mathbf{A}_t^T), \quad t = 1, 2, \dots \\ H_1 : \mathbf{y}_t | \mathbf{A}_t &\sim \mathcal{N}(0, \mathbf{A}_t \mathbf{A}_t^T), \quad t = 1, 2, \dots, \kappa, \\ &\mathbf{y}_t | \mathbf{A}_t \sim \mathcal{N}(\mathbf{A}_t \boldsymbol{\mu}, \mathbf{A}_t \mathbf{A}_t^T), \quad t = \kappa + 1, \kappa + 2, \dots \end{aligned} \quad (5)$$

Above, the distributions for the sketches are for given projections. Intuitively, for certain setting, the time-varying projection is preferred, e.g., when the post-change mean vector $\boldsymbol{\mu}$ is sparse, and the observations corresponding to missing data (i.e., only observe a subset of entries). One would expect observing a different subset of entries at each time would be better, because if the missing locations overlap with sparse mean shift locations, then we will miss the signal entirely.

2.2 Sketching matrices

In this paper, we assume that (1) when the sketching matrices \mathbf{A} or \mathbf{A}_t are random, then they have to be full row rank with probability 1; (2) when \mathbf{A} or \mathbf{A}_t are deterministic, then they have to be full row rank. The sketching matrices can either be user-specified or determined by the physical sensing system and not user specified. Below, we give several examples. Examples (i)–(iv) correspond to situations where the projections are user designed, and

example (v) (missing data) corresponds to the situation where the projections are imposed by the setup.

- (i) (Dimensionality reduction using random Gaussian matrices). To reduce the dimensionality of a high-dimensional vector (i.e., to compress data), one may use random projections. For instance, random Gaussian matrices $\mathbf{A} \in \mathbb{R}^{M \times N}$ whose entries are i.i.d. Gaussian with zero mean and variance equal to $1/M$.
- (ii) (Expander graphs). Sketching matrices with $\{0, 1\}$ entries are also commonly used: such a scenario is encountered in environmental monitoring (see, e.g., [15, 30]). Expander graphs are “sparse” 0-1 matrices in the sense that very few entries are zero and thus are desired for efficient computation since each linear projection only requires summing a few dimensions of the original data vector. Due to good structural properties, they have been used in compressed sensing (e.g., [31]). We will discuss more details about the expander graph in Section 4.4.3.
- (iii) (Pairwise comparison). In applications such as social network data analysis and computer vision, we are interested in a pairwise comparison of variables [32, 33]. This can be modeled as observing the difference between a pair of variables, i.e., at each time t , the measurements are $[\mathbf{x}_t]_i - [\mathbf{x}_t]_j$, for a set of $i \neq j$. There are a total of N^2 possible comparisons, and we may randomly select M out of N^2 such comparisons to observe. The pairwise comparison model leads to a structured fixed projection with only $\{0, 1, -1\}$ entries.
- (iv) (PCA). There are also approaches to change-point detection using principal component analysis (PCA) of the data streams (e.g., [28, 34]), which can be viewed as using a deterministic fixed projection \mathbf{A} , which is pre-computed as the signal singular space associated with the leading singular values of the data covariance matrix.
- (v) (Missing data). In various applications, we may only observe a subset of entries at each time (e.g., due to sensor failure), and the locations of the observed entries also vary with time [35]. This corresponds to $\mathbf{A}_t \in \mathbb{R}^{M_t \times N}$ being a submatrix of an identity matrix by selecting rows from an index set Ω_t at time t . When the data are missing at random, each entry of \mathbf{A}_t is i.i.d. Bernoulli random variables.

3 Methods: sketching procedures

Below, we derive the sketching procedure, when the projection matrices are fixed (across all times) and time-varying, respectively. In both cases, the MLE of the post-change mean vector cannot be uniquely determined generally. We tackle this issue and derive different generalized likelihood ratio (GLR) detection statistics and provide different analysis for the detection performances in the two cases.

3.1 Sketching procedure: fixed projection

3.1.1 Derivation of GLR statistic

We now derive the likelihood ratio statistic for the hypothesis test in (4). The strategy for deriving the GLR statistic in this case (with the fixed projection) is similar to [14]. However, [14] only considers the univariate case, where the MLE of the post-change mean can be obtained explicitly. Here, we consider the multi-dimensional case, and since the number of linear projections is much smaller than the original dimension, we cannot obtain a unique MLE for the post-change mean vector, but can only determine the equation that the MLE needs to satisfy; we need different derivation to obtain the GLR detection statistic.

Define the sample mean within a window $[k, t]$

$$\bar{\mathbf{y}}_{k,t} = \frac{\sum_{i=k+1}^t \mathbf{y}_i}{t-k}. \quad (6)$$

Since the observations are i.i.d. over time, for an assumed change-point $\kappa = k$, for the hypothesis test in (4), the log-likelihood ratio of observations accumulated up to time $t > k$, given the projection matrix \mathbf{A} , becomes

$$\begin{aligned} \ell(t, k, \boldsymbol{\mu}) &= \log \frac{\prod_{i=1}^k f_0(\mathbf{y}_i) \cdot \prod_{i=k+1}^t f_1(\mathbf{y}_i)}{\prod_{i=1}^t f_0(\mathbf{y}_i)} = \sum_{i=k+1}^t \log \frac{f_1(\mathbf{y}_i)}{f_0(\mathbf{y}_i)} \\ &= (t-k) \left[\bar{\mathbf{y}}_{k,t}^T (\mathbf{A}\mathbf{A}^T)^{-1} \mathbf{A}\boldsymbol{\mu} - \frac{1}{2} \boldsymbol{\mu}^T \mathbf{A}^T (\mathbf{A}\mathbf{A}^T)^{-1} \mathbf{A}\boldsymbol{\mu} \right], \end{aligned} \quad (7)$$

where $f_0 = \mathcal{N}(0, \mathbf{A}\mathbf{A}^T)$ denotes the probability density function of data under the null and $f_1 = \mathcal{N}(\mathbf{A}\boldsymbol{\mu}, \mathbf{A}\mathbf{A}^T)$ denotes the probability density function of \mathbf{y}_i under the alternative. Note that since \mathbf{A} is full row rank (with probability 1), $\mathbf{A}\mathbf{A}^T$ is invertible (with probability 1).

Since $\boldsymbol{\mu}$ is unknown, the GLR statistic replaces it with a maximum likelihood estimator (MLE) for fixed values of k and t in the likelihood ratio (7) to obtain the log-GLR statistic. Taking the derivative of $\ell(t, k, \boldsymbol{\mu})$ in (7) with respect to $\boldsymbol{\mu}$ and setting it to 0, we obtain an equation that the maximum likelihood estimator $\boldsymbol{\mu}^*$ of the post-change mean vector needs to satisfy:

$$\mathbf{A}^T (\mathbf{A}\mathbf{A}^T)^{-1} \mathbf{A}\boldsymbol{\mu}^* = \mathbf{A}^T (\mathbf{A}\mathbf{A}^T)^{-1} \bar{\mathbf{y}}_{k,t}, \quad (8)$$

or equivalently $\mathbf{A}^T [(\mathbf{A}\mathbf{A}^T)^{-1} \mathbf{A}\boldsymbol{\mu}^* - (\mathbf{A}\mathbf{A}^T)^{-1} \bar{\mathbf{y}}_{k,t}] = 0$. Note that since \mathbf{A}^T is of dimension M -by- N , this defines an under-determined system of equations for the maximum likelihood estimator $\boldsymbol{\mu}^*$. In other words, any $\boldsymbol{\mu}^*$ that satisfies

$$(\mathbf{A}\mathbf{A}^T)^{-1} \mathbf{A}\boldsymbol{\mu}^* = (\mathbf{A}\mathbf{A}^T)^{-1} \bar{\mathbf{y}}_{k,t} + \mathbf{c},$$

for a vector $\mathbf{c} \in \mathbb{R}^N$ that lies in the null space of \mathbf{A} , $\mathbf{A}^T \mathbf{c} = 0$, is a maximum likelihood estimator for the post-change mean. In this case, we could use pseudo-inverse to solve for $\boldsymbol{\mu}^*$, but we choose not to do this as the resulted

detection statistic is too complex to analyze. Rather, we choose a special solution by setting $\mathbf{c} = 0$, which will lead to a simple detection statistic and tractable theoretical analysis. Then, the corresponding maximum estimator satisfies the equation below:

$$(\mathbf{A}\mathbf{A}^T)^{-1} \mathbf{A}\boldsymbol{\mu}^* = (\mathbf{A}\mathbf{A}^T)^{-1} \bar{\mathbf{y}}_{k,t}. \quad (9)$$

Now substitute such a $\boldsymbol{\mu}^*$ into (7). Using (9), the first and second terms in (7) become, respectively,

$$\begin{aligned} \bar{\mathbf{y}}_{k,t}^T (\mathbf{A}\mathbf{A}^T)^{-1} \mathbf{A}\boldsymbol{\mu}^* &= \bar{\mathbf{y}}_{k,t}^T (\mathbf{A}\mathbf{A}^T)^{-1} \bar{\mathbf{y}}_{k,t}, \\ \frac{1}{2} \boldsymbol{\mu}^{*\top} \mathbf{A}^T (\mathbf{A}\mathbf{A}^T)^{-1} \mathbf{A}\boldsymbol{\mu}^* &= \frac{1}{2} \bar{\mathbf{y}}_{k,t}^T (\mathbf{A}\mathbf{A}^T)^{-1} \bar{\mathbf{y}}_{k,t}. \end{aligned}$$

Combining above, from (7), we have that the log-GLR statistic is given by

$$\ell(t, k, \boldsymbol{\mu}^*) = \frac{t-k}{2} \bar{\mathbf{y}}_{k,t}^T (\mathbf{A}\mathbf{A}^T)^{-1} \bar{\mathbf{y}}_{k,t}. \quad (10)$$

Since the change-point location k is unknown, when forming the detection statistic, we take the maximum over a set of possible locations, i.e., the most recent samples from $t-w$ to t , where $w > 0$ is the window size. Now we define the *sketching procedure*, which is a stopping time that stops whenever the log-GLR statistic raises above a threshold $b > 0$:

$$T = \inf \left\{ t : \max_{t-w \leq k < t} \frac{t-k}{2} \bar{\mathbf{y}}_{k,t}^T (\mathbf{A}\mathbf{A}^T)^{-1} \bar{\mathbf{y}}_{k,t} > b \right\}. \quad (11)$$

Here, the role of the window is twofold: it reduces the data storage when implementing the detection procedure and it establishes a minimum level of change that we want to detect.

3.1.2 Equivalent formulation of fixed projection sketching procedure

We can further simplify the log-GLR statistic in (10) using the singular value decomposition (SVD) of \mathbf{A} . This will facilitate the performance analysis in Section 4 and lead into some insights about the structure of the log-GLR statistic. Let the SVD of \mathbf{A} be given by

$$\mathbf{A} = \mathbf{U}\boldsymbol{\Sigma}\mathbf{V}^T, \quad (12)$$

where $\mathbf{U} \in \mathbb{R}^{M \times M}$, $\mathbf{V} \in \mathbb{R}^{N \times M}$ are the left and right singular spaces, $\boldsymbol{\Sigma} \in \mathbb{R}^{M \times M}$ is a diagonal matrix containing all non-zero singular values. Then $(\mathbf{A}\mathbf{A}^T)^{-1} = \mathbf{U}\boldsymbol{\Sigma}^{-2}\mathbf{V}^T$. Thus, we can write the log-GLR statistic (10) as

$$\ell(t, k, \boldsymbol{\mu}^*) = \frac{t-k}{2} \bar{\mathbf{y}}_{k,t}^T \mathbf{U}\boldsymbol{\Sigma}^{-2}\mathbf{U}^T \bar{\mathbf{y}}_{k,t}. \quad (13)$$

Substitution of the sample average (6) into (13) results in

$$\ell(t, k, \boldsymbol{\mu}^*) = \frac{\|\boldsymbol{\Sigma}^{-1}\mathbf{U}^T (\sum_{i=k+1}^t \mathbf{y}_i)\|^2}{2(t-k)}.$$

Now define transformed data

$$\mathbf{z}_i \triangleq \boldsymbol{\Sigma}^{-1}\mathbf{U}^T \mathbf{y}_i.$$

Since under the null hypothesis $\mathbf{y}_i|\mathbf{A} \sim \mathcal{N}(0, \mathbf{A}\mathbf{A}^\top)$, we have $\mathbf{z}_i \sim \mathcal{N}(0, \mathbf{I}_M)$. Similarly, under the alternative hypothesis $\mathbf{y}_i|\mathbf{A} \sim \mathcal{N}(\mathbf{A}\boldsymbol{\mu}, \mathbf{A}\mathbf{A}^\top)$, we have $\mathbf{z}_i \sim \mathcal{N}(\mathbf{V}^\top\boldsymbol{\mu}, \mathbf{I}_M)$. Combing above, we obtain the following equivalent form for the sketching procedure in (11):

$$T' = \inf \left\{ t : \max_{t-w \leq k < t} \frac{\|\sum_{i=k+1}^t \mathbf{z}_i\|^2}{2(t-k)} > b \right\}. \quad (14)$$

This form has one intuitive explanation: the sketching procedure essentially projects the data to form M (less than N) independent data streams and then form a log-GLR statistic for these independent data streams.

3.2 Sketching procedure: time-varying projection

3.2.1 GLR statistic

Similarly, we can derive the log likelihood ratio statistic for the time-varying projections. For an assumed change-point $\kappa = k$, using all observations from $k + 1$ to time t , we find the log likelihood ratio statistic similar to (7):

$$\begin{aligned} \ell(t, k, \boldsymbol{\mu}) &= \sum_{i=k+1}^t \left[\mathbf{y}_i^\top (\mathbf{A}_i \mathbf{A}_i^\top)^{-1} \mathbf{A}_i \boldsymbol{\mu} - \frac{1}{2} \boldsymbol{\mu}^\top \mathbf{A}_i^\top (\mathbf{A}_i \mathbf{A}_i^\top)^{-1} \mathbf{A}_i \mathbf{y}_i \right]. \end{aligned} \quad (15)$$

Similarly, we replace the unknown post-change mean vector $\boldsymbol{\mu}$ by its maximum likelihood estimator using data in $[k + 1, t]$. Taking the derivative of $\ell(t, k, \boldsymbol{\mu})$ in (15) with respect to $\boldsymbol{\mu}$ and setting it to 0, we obtain an equation that the maximum likelihood estimator $\boldsymbol{\mu}^*$ needs to satisfy

$$\left[\sum_{i=k+1}^t \mathbf{A}_i^\top (\mathbf{A}_i \mathbf{A}_i^\top)^{-1} \mathbf{A}_i \right] \boldsymbol{\mu}^* = \sum_{i=k+1}^t \mathbf{A}_i^\top (\mathbf{A}_i \mathbf{A}_i^\top)^{-1} \mathbf{y}_i. \quad (16)$$

Note that in the case of time-varying projection, we no longer have the structure in (8) for the fixed project. Thus, in this case, we will use a different strategy to derive the detection statistic based on pseudo-inverse. One needs to discuss the rank of the matrix $\sum_{i=k+1}^t \mathbf{A}_i^\top (\mathbf{A}_i \mathbf{A}_i^\top)^{-1} \mathbf{A}_i$ on the left-hand side of (16). Define the SVD of $\mathbf{A}_i = \mathbf{U}_i \mathbf{D}_i \mathbf{V}_i^\top$ with $\mathbf{U}_i \in \mathbb{R}^{M_i \times M_i}$ and $\mathbf{V}_i \in \mathbb{R}^{N \times M_i}$ being the eigenspaces and $\mathbf{D}_i \in \mathbb{R}^{M_i \times M_i}$ being a diagonal matrix that contains all the singular values. We have that

$$\sum_{i=k+1}^t \mathbf{A}_i^\top (\mathbf{A}_i \mathbf{A}_i^\top)^{-1} \mathbf{A}_i = \sum_{i=k+1}^t \mathbf{V}_i \mathbf{V}_i^\top = \mathbf{Q}\mathbf{Q}^\top, \quad (17)$$

where $\mathbf{Q} = [\mathbf{V}_{k+1}, \dots, \mathbf{V}_t] \in \mathbb{R}^{N \times S}$ and $S = \sum_{i=k+1}^t M_i$. Consider the rank of $\sum_{i=k+1}^t \mathbf{A}_i^\top (\mathbf{A}_i \mathbf{A}_i^\top)^{-1} \mathbf{A}_i$ for two cases below:

From above, we can see that this matrix is rank deficient when $t - k < N/M$, i.e., the number of post-change samples $t - k$ is small. However, this is generally the case since

we want to detect the change quickly once it occurs. Since the matrix in (17) is non-invertible in general, we use the pseudo-inverse of the matrix. Define

$$\mathbf{B}_{k,t} \triangleq \left(\sum_{i=k+1}^t \mathbf{A}_i^\top (\mathbf{A}_i \mathbf{A}_i^\top)^{-1} \mathbf{A}_i \right)^\dagger \in \mathbb{R}^{N \times N}.$$

From (16), we obtain an estimate of the maximum likelihood estimator for the post-change mean

$$\boldsymbol{\mu}^* = \mathbf{B}_{k,t} \sum_{i=k+1}^t \mathbf{A}_i^\top (\mathbf{A}_i \mathbf{A}_i^\top)^{-1} \mathbf{y}_i.$$

Substituting such a $\boldsymbol{\mu}^*$ into (15), we obtain the log-GLR statistic for time-varying projection:

$$\begin{aligned} \ell(t, k, \boldsymbol{\mu}^*) &= \frac{1}{2} \left(\sum_{i=k+1}^t \mathbf{A}_i^\top (\mathbf{A}_i \mathbf{A}_i^\top)^{-1} \mathbf{y}_i \right)^\top \mathbf{B}_{k,t} \\ &\quad \cdot \left(\sum_{i=k+1}^t \mathbf{A}_i^\top (\mathbf{A}_i \mathbf{A}_i^\top)^{-1} \mathbf{y}_i \right). \end{aligned} \quad (18)$$

3.2.2 Time-varying 0-1 project matrices

To further simplify the expression of GLR in (18), we consider a special case when \mathbf{A}_t has only one entry equal to 1 for each row and all other entries equal to 0. This means that at each time, we only observe a subset of the entries and can correspond to the missing data case. Now $\mathbf{A}_t \mathbf{A}_t^\top$ is an M_t -by- M_t identity matrix, and $\mathbf{A}_t^\top \mathbf{A}_t$ is a diagonal matrix. For a diagonal matrix $\mathbf{D} \in \mathbb{R}^{N \times N}$ with diagonal entries $\lambda_1, \dots, \lambda_N$, the pseudo-inverse of \mathbf{D} is a diagonal matrix with diagonal entries λ_i^{-1} if $\lambda_i \neq 0$ and with diagonal entries 0 if $\lambda_i = 0$. Let the index set of the observed entries at time t be Ω_t . Define indicator variables

$$\mathbb{I}_{tn} = \begin{cases} 1, & \text{if } n \in \Omega_t; \\ 0, & \text{otherwise.} \end{cases} \quad (19)$$

Then, the log-GLR statistic in (18) becomes

$$\ell(t, k, \boldsymbol{\mu}^*) = \sum_{n=1}^N \frac{(\sum_{i=k+1}^t [\mathbf{x}_i]_n \mathbb{I}_{in})^2}{\sum_{i=k+1}^t \mathbb{I}_{in}}, \quad (20)$$

Hence, for 0-1 matrices, the sketching procedure based on log-GLR statistic is given by

$$T_{\{0,1\}} = \inf \left\{ t : \max_{t-w \leq k < t} \frac{1}{2} \sum_{n=1}^N \frac{(\sum_{i=k+1}^t [\mathbf{x}_i]_n \mathbb{I}_{in})^2}{\sum_{i=k+1}^t \mathbb{I}_{in}} > b \right\}, \quad (21)$$

where $b > 0$ is the prescribed threshold and w is the window length. Note that the log-GLR statistic essentially computes the sum of each entry within the time window $[t - w, t)$ and then averages the squared sum.

4 Results: Theoretical

In this section, we present theoretical approximations to two performance metrics, the average run length (ARL), which captures the false alarm rate, and the expected detection delay (EDD), which captures the power of the detection statistic.

4.1 Performance metrics

We first introduce some necessary notations. Under the null hypothesis in (1), the observations are zero mean. Denote the probability and expectation in this case by \mathbb{P}^∞ and \mathbb{E}^∞ , respectively. Under the alternative hypothesis, there exists a change-point κ , $0 \leq \kappa < \infty$ such that the observations have mean $\boldsymbol{\mu}$ for all $t > \kappa$. Probability and expectation in this case are denoted by \mathbb{P}^κ and \mathbb{E}^κ , respectively.

The choice of the threshold b involves a tradeoff between two standard performance metrics that are commonly used for analyzing change-point detection procedures [15]: (i) the ARL, defined to be the expected value of the stopping time when there is no change, and (ii) the EDD, defined to be the expected stopping time in the extreme case where a change occurs immediately at $\kappa = 0$, which is denoted as $\mathbb{E}^0\{T\}$.

The following argument from [14] explains why we consider $\mathbb{E}^0\{T\}$. When there is a change at κ , we are interested in the expected delay until its detection, i.e., the conditional expectation $\mathbb{E}^\kappa\{T - \kappa | T > \kappa\}$, which is a function of κ . When the shift in the mean only occurs in the positive direction $[\boldsymbol{\mu}]_i \geq 0$, it can be shown that $\sup_\kappa \mathbb{E}^\kappa\{T - \kappa | T > \kappa\} = \mathbb{E}^0\{T\}$. It is not obvious that this remains true when $[\boldsymbol{\mu}]_i$ can be either positive or negative. However, since $\mathbb{E}^0\{T\}$ is certainly of interest and reasonably easy to analyze, it is common to consider $\mathbb{E}^0\{T\}$ in the literature and we also adopt this as a surrogate.

4.2 Fixed projection

Define a special function (cf. [36], page 82)

$$\nu(u) = 2u^{-2} \exp \left[-2 \sum_{i=1}^{\infty} i^{-1} \Phi(-|u|i^{1/2}/2) \right],$$

where Φ denotes the cumulative probability function (CDF) for the standard Gaussian with zero mean and unit variance. For numerical purposes, a simple and accurate approximation is given by (cf. [37])

$$\nu(u) \approx \frac{2/u[\Phi(u/2) - 0.5]}{(u/2)\Phi(u/2) + \phi(u/2)},$$

where ϕ denotes the probability distribution function (PDF) for standard Gaussian. We obtain an approximation to the ARL of the sketching procedure with a fixed projection as follows:

Theorem 1 (ARL, fixed projection) *Assume that $1 \leq M \leq N$, $b \rightarrow \infty$ with $M \rightarrow \infty$ and b/M fixed. Then, with $w = o(b^r)$ for some positive integer r , for a given projection matrix \mathbf{A} that is full rank deterministically or with probability 1, the ARL of the sketching procedure defined in (11) is given by*

$$\mathbb{E}^\infty\{T\} = \frac{2\sqrt{\pi}}{c(M, b, w)} \frac{1}{1 - \frac{M}{2b}} \frac{1}{\sqrt{M}} \left(\frac{M}{2b}\right)^{\frac{M}{2}} e^{b - \frac{M}{2}} (1 + o(1)), \quad (22)$$

where

$$c(M, b, w) = \int_{\sqrt{\frac{2b}{w}(1 - \frac{M}{2b})}}^{\sqrt{2b(1 - \frac{M}{2b})}} uv^2(u) du. \quad (23)$$

This theorem gives an explicit expression for ARL as a function of the threshold b , the dimension of the sketches M , and the window length w . As we will show below, the approximation to ARL given by Theorem 1 is highly accurate. On a higher level, this theorem characterizes the mean of the stopping time, when the detection statistic is driven by noise. The requirement for $w = o(b^r)$ for some positive integer r comes from [15] that our results are based on; this ensures the correct scaling when we pass to the limit. This essentially requires that the window length be large enough when the threshold b increases. In practice, w has to be large enough so that it does not cause a miss detection: w has to be longer than the anticipated expected detection delay as explained in [15].

Moreover, we obtain an approximation to the EDD of the sketching procedure with a fixed projection as follows. Define

$$\Delta = \|\mathbf{V}^\top \boldsymbol{\mu}\|, \quad (24)$$

where \mathbf{V} contains the left singular vectors of \mathbf{A} . Let $\tilde{S}_t \triangleq \sum_{i=1}^t \delta_i$ be a random walk where the increments δ_i are independent and identically Gaussian distributed with mean $\Delta^2/2$ and variance Δ^2 . We can find the expected value of the minimum is given by [37]

$$\mathbb{E} \left\{ \min_{t \geq 0} \tilde{S}_t \right\} = - \sum_{i=1}^{\infty} i^{-1} \mathbb{E} \left\{ \tilde{S}_i^- \right\},$$

where $(x)^- = -\min\{x, 0\}$, and the infinite series converges and can be evaluated easily numerically. Also, define

$$\rho(\Delta) = \Delta^2/4 + 1 - \sum_{i=1}^{\infty} i^{-1} \mathbb{E} \left\{ \tilde{S}_i^- \right\}.$$

Theorem 2 (EDD, fixed projection) *Suppose $b \rightarrow \infty$ with other parameters held fixed. Then, for a given matrix \mathbf{A} with the right singular vectors \mathbf{V} , the EDD of the sketching procedure (11) when $\kappa = 0$ is given by*

$$\mathbb{E}^0\{T\} = \frac{b + \rho(\Delta) - M/2 + \mathbb{E}\{\min_{t \geq 0} \tilde{S}_t\} + o(1)}{\Delta^2/2}, \quad (25)$$

The theorem finds an explicit expression for the EDD as a function of threshold b , the number of sketches M , and the signal strength captured through Δ which depends on the post mean vector μ and the projection subspace V .

The proofs for the above two theorems utilize the equivalent form of T in (14) and draw a connection of the sketching procedure to the so-called mixture procedure (cf. T_2 in [15]) when M sensors are affected by the change, and the post-change mean vector is given by $V^T \mu$.

4.2.1 Accuracy of theoretical approximations

Consider a A generated as a Gaussian random matrix, with entries i.i.d. $\mathcal{N}(0, 1/N)$. Using the expression in Theorem 1, we can find the threshold b such that the corresponding ARL is equal to 5000. This can be done conveniently; since the ARL is an increasing function of the threshold b , we can use bisection to find such a threshold b . Then, we compare it with a threshold b found from the simulation.

As shown in Table 1, the threshold found using Theorem 1 is very close to that obtained from simulation. Therefore, even if the theoretical ARL approximation is derived for N tends to infinity, it is still applicable when N is large but finite. Theorem 1 is quite useful in determining a threshold for a targeted ARL, as simulations for large N and M can be quite time-consuming, especially for a large ARL (e.g., 5000 or 10,000).

Moreover, we simulate the EDD for detecting a signal such that the post-change mean vector μ has all entries equal to a constant $[\mu]_i = 0.5$. As also shown in Table 1, the approximation for EDD using Theorem 2 is quite accurate.

We have also verified that the theoretical approximations are accurate for the expander graphs and details omitted here since they are similar.

4.2.2 Consequence

Theorems 1 and 2 have the following consequences:

Table 1 Verification of numerical accuracy of theoretical results

M	b (theo)	b (simu)	EDD (theo)	EDD (simu)
100	84.65	84.44	3.4	4.3 (0.9)
70	64.85	64.52	4.0	5.1 (1.2)
50	51.04	50.75	4.8	5.9 (1.6)
30	36.36	36.43	7.7	7.6 (2.5)
10	19.59	19.63	19.8	17.4 (9.8)

A being a fixed Gaussian random matrix. $N = 100$, $w = 200$, $ARL = 5000$, for simulated EDD $[\mu]_i = 0.5$. Numbers in the parentheses are the standard deviation of the simulated EDD

Remarks 1 For a fixed large ARL, when M increases, the ratio M/b is bounded between 0.5 and 2. This is a property quite useful for establishing results in Section 4.4. This is demonstrated numerically in Fig. 1 when $N = 100$, $w = 200$, for a fixed ARL being 5000. The corresponding threshold b is found using Theorem 1, when M increases from 10 to 100. More precisely, Theorem 1 leads to the following corollary:

Corollary 1 Assume a large constant $\gamma \in (e^5, e^{20})$. Let $w \geq 100$. For any large enough $M > 24.85$, the threshold b such that the corresponding ARL $\mathbb{E}^\infty\{T\} = \gamma$ satisfies $M/b \in (0.5, 2)$. In other words, $\max\{M/b, b/M\} \leq 2$.

Note that e^{20} is on the order of 5×10^8 ; hence, it means that ARL can be very large; however, it is still bounded above (this means that the corollary holds for a non-asymptotic regime).

Remarks 2 As $b \rightarrow \infty$, the first order approximation to the EDD in Theorem 2 is given by $b/(\Delta^2/2)$, i.e., the threshold b divided by the Kullback-Leibler (K-L) divergence (see, e.g., [15] shows that $\Delta^2/2$ is the K-L divergence between $\mathcal{N}(0, I_M)$ and $\mathcal{N}(V^T \mu, I_M)$). This is consistent with our intuition since the expected increment of the detection statistics is roughly the K-L divergence of the test. For finite b , especially when the signal strength is weak and when the number of sketches M is not large enough, the other terms than $b/(\Delta^2/2)$ will play a significant role in determining the EDD.

4.3 Time-varying 0-1 random projection matrices

Below, we obtain approximations to ARL and EDD for $T_{\{0,1\}}$, i.e., when 0-1 sketching matrices are used. We

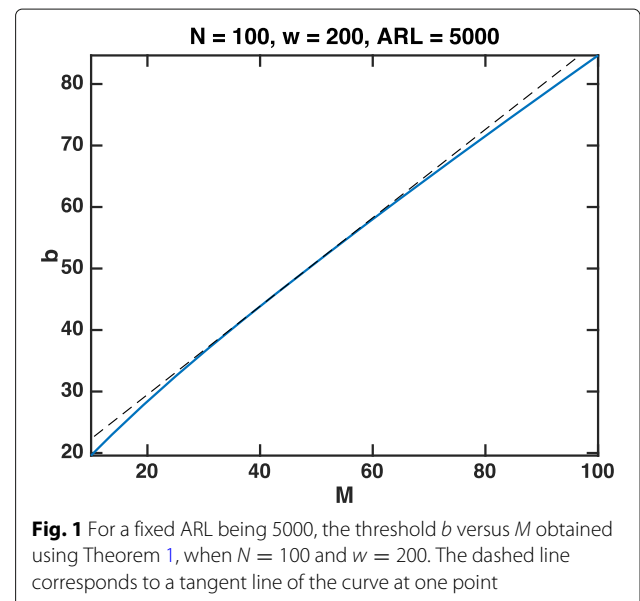


Fig. 1 For a fixed ARL being 5000, the threshold b versus M obtained using Theorem 1, when $N = 100$ and $w = 200$. The dashed line corresponds to a tangent line of the curve at one point

assume a fixed dimension $M_t = M, \forall t > 0$. We also assume that at each time t , we randomly select M out of N dimensions to observe. Hence, at each time, each signal dimension has a probability

$$r = M/N \in (0, 1)$$

to be observed. The sampling scheme is illustrated in Fig. 2, when $N = 10$ and $M = 3$ (the number of the dots in each column is 3) over 17 consecutive time periods from time $k = t - 17$ to time t .

For such a sampling scheme, we have the following result:

Theorem 3 (ARL, time-varying 0-1 random projection)

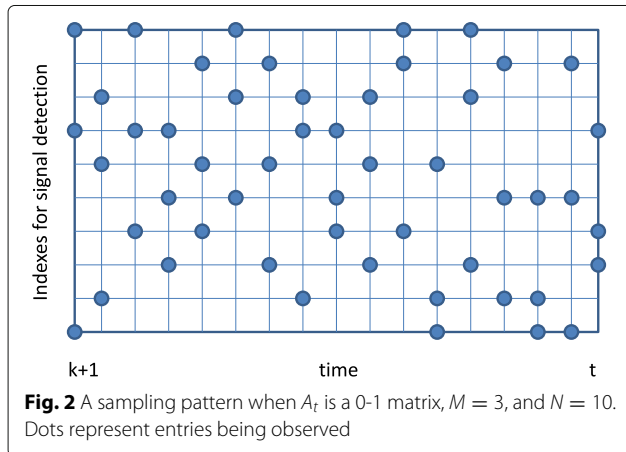
Let $r = M/N$. Let $b' = b/r$. When $b \rightarrow \infty$, for the procedure defined in (21), we have that

$$\begin{aligned} \mathbb{E}^\infty \{T_{\{0,1\}}\} \\ = \frac{2\sqrt{\pi}}{c(N, b', w)} \frac{1}{\sqrt{N}} \frac{1}{1 - \frac{N}{2b'}} \left(\frac{N}{2}\right)^{\frac{N}{2}} b'^{-\frac{N}{2}} e^{b' - \frac{N}{2}} + o(1), \end{aligned} \quad (26)$$

where $c(N, b', w)$ is defined by replacing b with b' in (23).

Moreover, we can obtain an approximation to EDD of $T_{\{0,1\}}$, as justified by the following arguments. First, relax the deterministic constraint that at each time we observe exactly M out of N entries. Instead, assume a random sampling scheme that at each time we observe one entry of x_i with probability r , $1 \leq n \leq N$. Consider i.i.d. Bernoulli random variables ξ_{ni} with parameter r for $1 \leq n \leq N$ and $i \geq 1$. Define

$$Z_{n,k,t} \triangleq \frac{\sum_{i=k+1}^t [x_i]_n \xi_{ni}}{\sqrt{(t-k)r}}.$$



Based on this, we define a procedure whose behavior is arguably similar to $T_{\{0,1\}}$:

$$T'_{\{0,1\}} = \inf \left\{ t \geq 1 : \max_{t-w \leq k < t} \frac{1}{2} \sum_{n=1}^N Z_{n,k,t}^2 > b \right\},$$

where $b > 0$ is the prescribed threshold. Then, using the arguments in Appendix 2, we can show that the approximation to EDD of this procedure is given by

$$\mathbb{E}^0 \{T'_{\{0,1\}}\} = \left(\frac{2b - N}{\sum_{n=1}^N \mu_n^2} + o(1) \right) \cdot \frac{N}{M}, \quad (27)$$

and we use this to approximate the EDD of $T_{\{0,1\}}$.

Table 2 shows the accuracy of the approximations for ARL in (26) and for EDD in (27) with various M 's when $N = 100$, $w = 200$, and all entries of $[\mu]_i = 0.5$. The results show that the thresholds b obtained using the theoretical approximations and that the EDD approximations are both very accurate.

4.4 Bounding relative performance

In this section, we characterize the relative performance of the sketching procedure compared to that without sketching (i.e., using the original log-GLR statistic). We show that the performance loss due sketching can be small, when the signal-to-noise ratio and M are both sufficiently large. In the following, we focus on fixed projection to illustrate this point.

4.4.1 Relative performance metric

We consider a relative performance measure, which is the ratio of EDD using the original data (denoted as $\text{EDD}(N)$), which corresponds to $A = I$, versus the EDD using the sketches (denoted as $\text{EDD}(M)$)

$$\frac{\text{EDD}(N)}{\text{EDD}(M)} \in (0, 1).$$

We will show that this ratio depends critically on the following quantity

$$\Gamma \triangleq \frac{\|V^T \mu\|^2}{\|\mu\|^2}, \quad (28)$$

Table 2 Verification of numerical accuracy of theoretical results

M	b (theo)	b (simu)	EDD (theo)	EDD (simu)
100	84.65	84.44	2.8	3.3 (0.8)
70	83.72	83.41	3.8	4.5 (1.2)
50	82.84	83.02	5.3	6.1 (1.5)
30	81.46	82.48	8.7	9.8 (2.4)
10	78.32	79.27	23.4	26.6 (6.4)

A_t 's being time-varying. $N = 100$, $w = 200$, ARL = 5000; for simulated EDDs, all $[\mu]_i = 0.5$. Numbers in the parentheses are standard deviation of the simulated results

which is the ratio of the KL divergence after and before the sketching.

We start by deriving the relative performance measure using theoretical approximations we obtained in the last section. Recall the expression for EDD approximation in (25). Define

$$h(\Delta, M) = \rho(\Delta) - M/2 - \mathbb{E}\{\tilde{S}_i^-\}. \tag{29}$$

From Theorem 2, we obtain that the EDD of the sketching procedure is proportional to

$$\frac{2b}{\|V^T \mu\|^2} \cdot \left(1 + \frac{h(\|V^T \mu\|, M)}{2b}\right) \cdot (1 + o(1)).$$

Let b_N and b_M be the thresholds such that the corresponding ARLs are 5000, for the procedure without sketching and with M sketches, respectively. Define $Q_M = M/b_M$, $Q_N = N/b_N$ and

$$P = \frac{1 + h(\|\mu\|, N)/b_N}{1 + h(\|V^T \mu\|, M)/b_M}. \tag{30}$$

Using the definitions above, we have

$$\begin{aligned} \frac{\text{EDD}(N)}{\text{EDD}(M)} &= P \cdot \frac{b_N}{b_M} \cdot \frac{\|V^T \mu\|^2}{\|\mu\|^2} (1 + o(1)) \\ &= P \cdot \frac{N}{M} \cdot \frac{Q_M}{Q_N} \cdot \Gamma(1 + o(1)). \end{aligned} \tag{31}$$

4.4.2 Discussion of factors in (31)

We can show that $P \geq 1$ for sufficiently large M and large signal strength. This can be verified numerically. Since all quantities that P depends on can be computed explicitly: the thresholds b_N and b_M can be found from Theorem 1 once we set a target ARL, the h function can be evaluated using (29) which depends explicitly on Δ and M . Figure 3 shows the value of P when $N = 100$ and all the entries of the post-change mean vector $[\mu]_i$ are equal to a constant value that varies across the x -axis. Note that P is less than 1 only when the signal strength μ_i are small and M is small. Thus, we have,

$$\frac{\text{EDD}(N)}{\text{EDD}(M)} \geq \frac{N}{M} \cdot \frac{Q_M}{Q_N} \cdot \Gamma(1 + o(1)),$$

for sufficiently large M and signal strength Δ .

Using Corollary 1, we have that $Q_M \in (0.5, 2)$ and $Q_N \in (0.5, 2)$, and hence, a lower bound of the ratio $\text{EDD}(N)/\text{EDD}(M)$ is between $(1/4)(N/M)\Gamma$ and $4(N/M)\Gamma$, for large M or large signal strength.

Next, we will bound Γ when A is a Gaussian matrix and an expander graph, respectively.

4.4.3 Bounding Γ

Gaussian matrix. Consider $A \in \mathbb{R}^{M \times N}$ whose entries are i.i.d. Gaussian with zero mean and variance equal to $1/M$. First, we have the following lemma

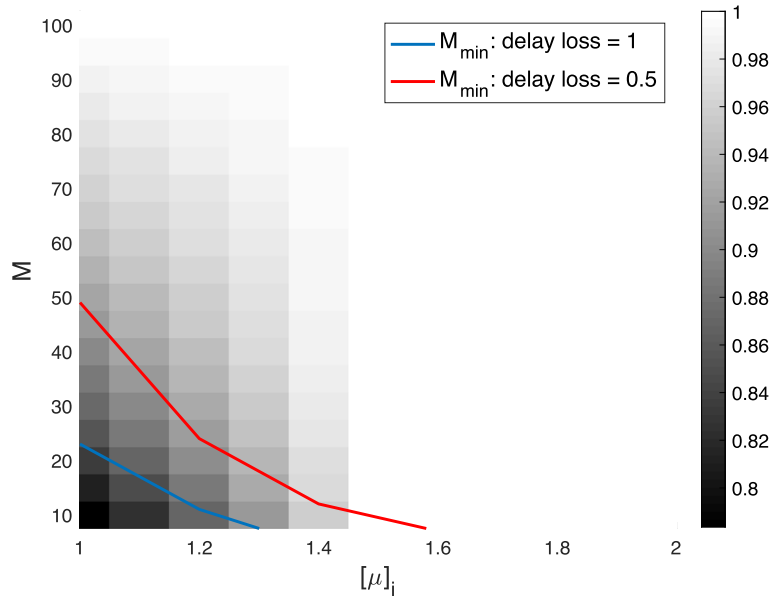


Fig. 3 The P factor defined in (30) for different M and $[\mu]_i$, when the post-change mean vector has entries all equal to $[\mu]_i$. Assume $N = 100$. The white regions correspond to $P \geq 1$, and dark regions correspond to $P < 1$ and the darker, the smaller P is (note that the smallest P in this graph is above 0.75). We also plot the M_{\min} (defined later in (35)) required in these cases such that the EDD of the sketching procedure is no more than δ larger than the corresponding procedure without sketching (fixing ARL = 5000), for $\delta = 1$ and $\delta = 0.5$. The M_{\min} are obtained by Monte Carlo simulation. The M_{\min} versus $[\mu]_i$ correspond to the blue and the red curves, respectively. Above these two curves, the EDD with sketching is almost the same as before (without sketching), i.e., the regime where sketching has little loss. The left-bottom corner corresponds to the region where sketching has more loss. This also shows that indeed $P < 1$ is an indicator of significant performance loss using sketching

Lemma 1 ([38]) *Let $A \in \mathbb{R}^{M \times N}$ have i.i.d. $\mathcal{N}(0, 1)$ entries. Then, for any fixed vector μ , we have that*

$$\Gamma \sim \text{Beta}\left(\frac{M}{2}, \frac{N-M}{2}\right). \tag{32}$$

More related results can be found in [39]. Since the $\text{Beta}(\alpha, \beta)$ distribution has a mean $\alpha/(\alpha + \beta)$, we have that

$$\mathbb{E}\{\Gamma\} = \frac{M/2}{M/2 + (N-M)/2} = \frac{M}{N}.$$

We may also show that, provided M and N grow proportionally, Γ converges to its mean value at a rate exponential in N . Define $\delta \in (0, 1)$ to be

$$\delta \triangleq \lim_{N \rightarrow \infty} \frac{M}{N}. \tag{33}$$

We have the following result.

Theorem 4 (Gaussian A) *Let $A \in \mathbb{R}^{M \times N}$ have entries i.i.d. $\mathcal{N}(0, 1)$. Let $N \rightarrow \infty$ such that (33) holds. Then, for $0 < \epsilon < \min(\delta, 1 - \delta)$, we have that*

$$\mathbb{P}\{\delta - \epsilon < \Gamma < \delta + \epsilon\} \rightarrow 1. \tag{34}$$

Hence, for Gaussian A , Γ is approximately M/N with probability 1.

Note that Theorem 4 is different from the restricted isometry property (RIP) invoked in compressed sensing, since here we assume one fixed and given vector μ , but in compressed sensing, one cares about norm preservation uniformly for all sparse vectors (with the same sparsity level) with probability 1.

Expander graph A . We can show that for expander graphs, Γ is also bounded. This holds for the “one-sided” changes, i.e., the post-change mean vector is element-wise positive.

A matrix A corresponds to a (s, ϵ) -expander graph with regular right degree d if and only if each column of A has exactly d “1”s, and for any set S of right nodes with $|S| \leq s$, the set of neighbors $\mathcal{N}(S)$ of the left nodes has size $|\mathcal{N}(S)| \geq (1 - \epsilon)d|S|$. If it further holds that each row of A has c “1”s, we say A corresponds to a (s, ϵ) -expander with regular right degree d and regular left degree c .

Assume $[\mu]_i \geq 0$ for all i . Let $A \in \mathbb{R}^{M \times N}$ be consisting of binary entries, which corresponds to a bipartite graph, illustrated in Fig. 4. We further consider a bipartite graph with regular left degree c (i.e., the number of edges from each variable node is c) and regular right degree d (i.e., the number of edges from each parity check node is d), as illustrated in Fig. 4. Hence, this requires $Nc = Md$. Expander graphs satisfy the above requirements. The existence of expander graphs is established in [40]:

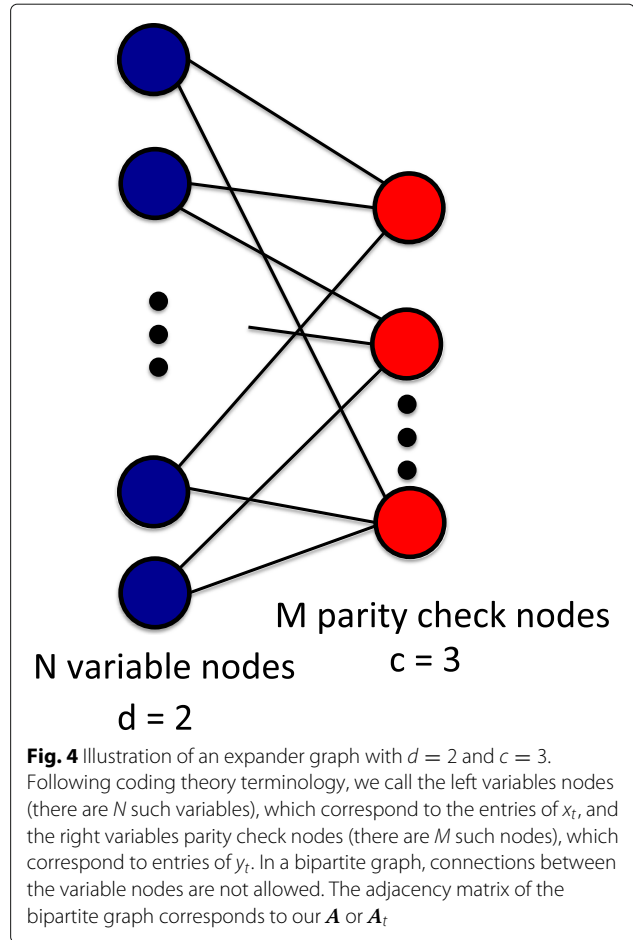


Fig. 4 Illustration of an expander graph with $d = 2$ and $c = 3$. Following coding theory terminology, we call the left variables nodes (there are N such variables), which correspond to the entries of x_t , and the right variables parity check nodes (there are M such nodes), which correspond to entries of y_t . In a bipartite graph, connections between the variable nodes are not allowed. The adjacency matrix of the bipartite graph corresponds to our A or A_t

Lemma 2 ([40]) *For any fixed $\epsilon > 0$ and $\rho \triangleq M/N < 1$, when N is sufficiently large, there always exists an $(\alpha N, \epsilon)$ expander with a regular right degree d and a regular left degree c for some constants $\alpha \in (0, 1)$, d and c .*

Theorem 5 (Expander A) *If A corresponds to a (s, ϵ) -expander with regular degree d and regular left degree c , for any nonnegative vector $[\mu]_i \geq 0$, we have that*

$$\Gamma \geq \frac{M(1 - \epsilon)}{dN}.$$

Hence, for expander graphs, Γ is approximately greater than $M/N \cdot (1/d)$, where d is a small number.

4.4.4 Consequence

Combine the results above, we showed that for the regime where M and the signal strength are sufficiently large, the performance loss can be small (as indeed observed from our numerical examples). In this regime, when A is a Gaussian random matrix, the relative performance measure $\text{EDD}(N)/\text{EDD}(M)$ is a constant, under the conditions in Corollary 1. Moreover, when A is a sparse 0-1 matrix with d non-zero entries on each row (in particular,

an expander graph), the ratio (31) $EDD(N)/EDD(M)$ is lower bounded by $(1/4) \cdot d/(1 - \epsilon)$ for some small number $\epsilon > 0$, when Corollary 1 holds.

There is one intuitive explanation. Unlike in compressed sensing, where the goal is to recover a sparse signal and one needs the projection to preserve norm up to a factor through the restricted isometry property (RIP) [41], our goal is to detect a non-zero vector in Gaussian noise, which is a much simpler task than compressed sensing. Hence, even though the projection reduces the norm of the vector, as long as the projection does not diminish the signal is normal below the noise floor.

On the other hand, when the signal is weak, and M is not large enough, there can be significant performance loss (as indeed observed in our numerical examples) and we cannot lower bound the relative performance measure. Fortunately, in this regime, we can use our theoretical results in Theorems 1 and 2 to design the number of sketches M for an anticipated worst-case signal strength Δ , or determine the infeasibility of the problem, i.e., it is better not to use sketching since the signal is too weak.

5 Results: numerical examples

In this section, we present numerical examples to demonstrate the performance of the sketching procedure. We focus on comparing the sketching procedure with the GLR procedure without sketching (by letting $A = I$ in the sketching procedure). We also compare the sketching procedures with a standard multivariate CUSUM using sketches.

In the subsequent examples, we select ARL to be 5000 to represent a low false detection rate (similar choice has been made in other sequential change-point detection work such as [15]). In practice, however, the target ARL value depends on how frequent we can tolerate false detection (e.g., once a month or once a year). Below, EDD_o denotes the EDD when $A = I$ (i.e., no sketching is used). All simulated results are obtained from 10^4 repetitions. We also consider the minimum number of sketches

$$M_{\min} = \arg \min\{M : EDD(M) \leq \delta + EDD_o\}, \quad (35)$$

such that the corresponding sketching procedure is only δ sample slower than the full procedure. Below, we focus on the delay loss $\delta = 1$.

5.1 Fixed projection, Gaussian random matrix

First, consider Gaussian A with $N = 500$ and different number of sketches $M < N$.

5.1.1 EDD versus signal magnitude

Assume the post-change mean vector has entries with equal magnitude: $[\mu]_i = \mu_0$, to simplify our discussion. Figure 5a shows EDD versus an increasing signal magnitude μ_0 . Note that when μ_0 and M are sufficiently large,

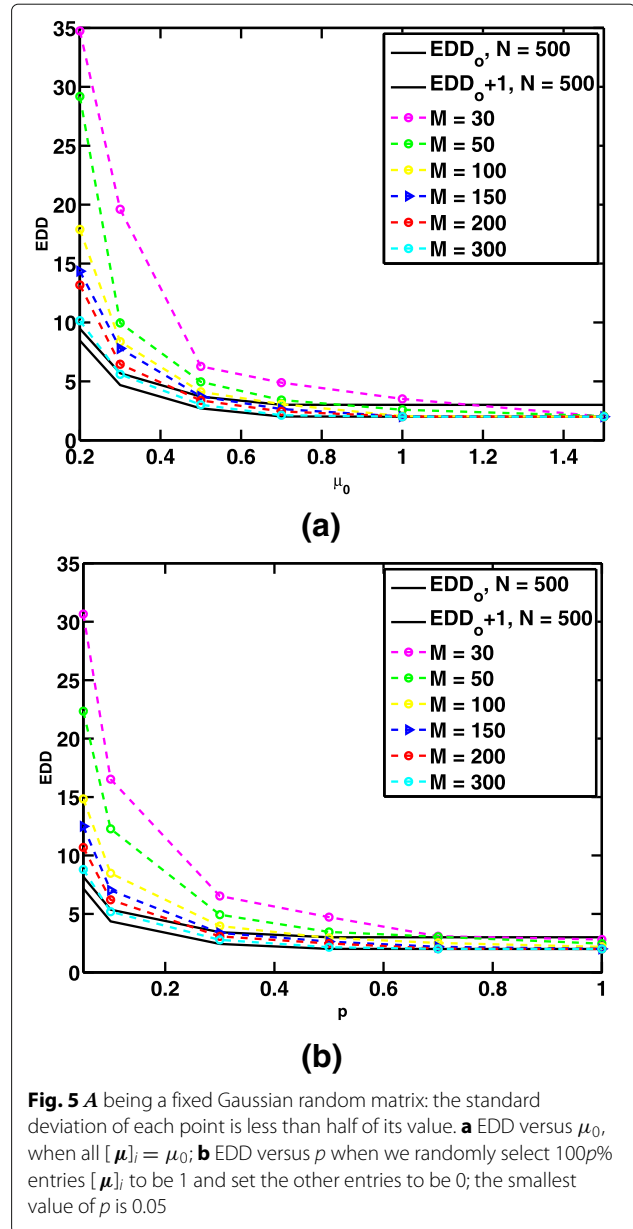


Fig. 5 A being a fixed Gaussian random matrix: the standard deviation of each point is less than half of its value. **a** EDD versus μ_0 , when all $[\mu]_i = \mu_0$; **b** EDD versus p when we randomly select $100p\%$ entries $[\mu]_i$ to be 1 and set the other entries to be 0; the smallest value of p is 0.05

the sketching procedure can approach the performance of the procedure using the full data as predicted by our theory. When signal is weak, we have to use a much larger M to prevent a significant performance loss (and when signal is too weak, we cannot use sketching). Table 3 shows M_{\min} for each signal strength; we find that when μ_0 is sufficiently large, we may even use M_{\min} less than 30 for an

Table 3 Assume A being a fixed Gaussian random matrix

μ_0	0.3	0.5	0.7	1	1.2
M_{\min}	300	150	100	50	30

The table shows M_{\min} required for various mean shifts μ_0 as shown in Fig. 5a. Here $N = 500$, $w = 200$, and all $[\mu]_i = \mu_0$

$N = 500$ to have little performance loss. Note that here, we do not require signals to be sparse.

5.1.2 EDD versus signal sparsity

Now assume that the post-change mean vector is sparse: only $100p\%$ entries μ_i being 1 and the other entries being 0. Figure 5b shows EDD versus an increasing p . Note that as p increases, the signal strength also increases; thus, the sketching procedure will approach the performance using the full data. Similarly, the M_{\min} required is listed in Table 4. For example, when $p = 0.5$, we find that one can use $M_{\min} = 100$ for an $N = 500$ with little performance loss.

5.2 Fixed projection, expander graph

Now assume A being an expander graph with $N = 500$ and different number of sketches $M < N$. We run the simulations with the same settings as those in Section 5.1.

5.2.1 EDD versus signal magnitude

Assume the post-change mean vector $[\mu]_i = \mu_0$. Figure 6a shows EDD with an increasing μ_0 . Note that the simulated EDDs are smaller than those for the Gaussian random projections in Fig. 5. A possible reason is that the expander graph is better at aggregating the signals when $[\mu]_i$ are all positive. However, when $[\mu]_i$ can be either positive or negative, the two choices of A have similar performance, as shown in Fig. 7, where $[\mu]_i$ are drawn i.i.d. uniformly from $[-3, 3]$.

5.2.2 EDD versus signal sparsity

Assume that the post-change mean vector has only $100p\%$ entries $[\mu]_i$ being 1 and the other entries being 0. Figure 6b shows the simulated EDD versus an increasing p . As p tends to 1, the sketching procedure approaches the performance using the full data.

5.3 Time-varying projections with 0-1 matrices

To demonstrate the performance of the procedure $T_{\{0,1\}}$ (21) using time-varying projection with 0-1 entries, again, we consider two cases: the post-change mean vector $[\mu]_i = \mu_0$ and the post-change mean vector has $100p\%$ entries $[\mu]_i$ being 1 and the other entries being 0. The simulated EDDs are shown in Fig. 8. Note that $T_{\{0,1\}}$ can detect change quickly with a small subset of observations. Although EDDs of $T_{\{0,1\}}$ are larger than those for the fixed

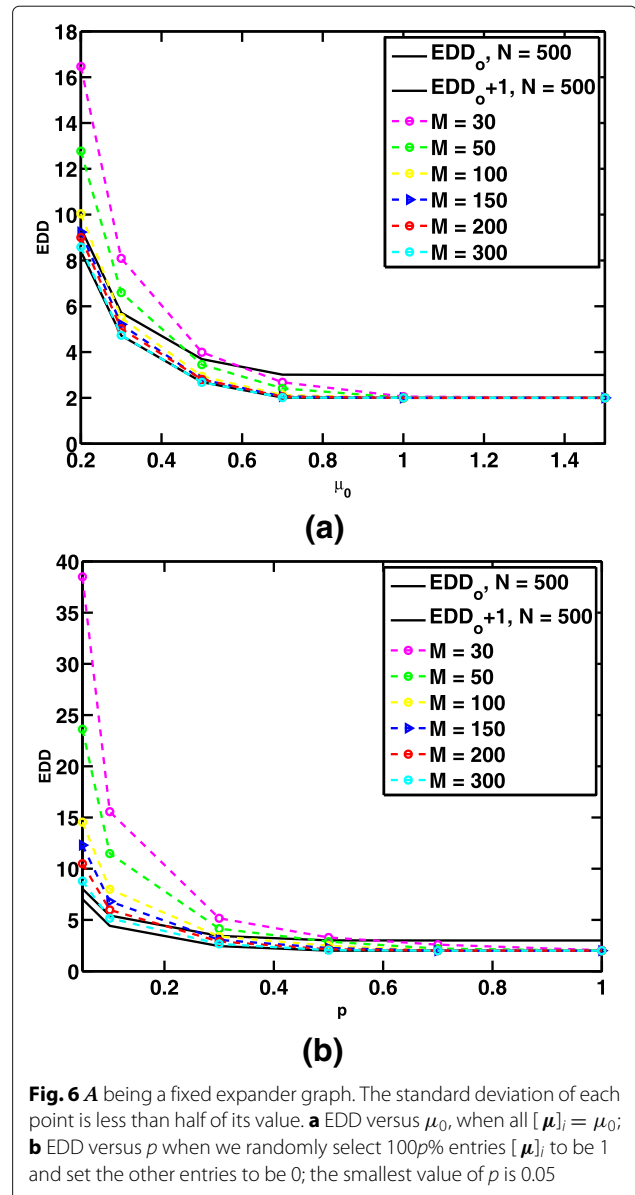


Fig. 6 A being a fixed expander graph. The standard deviation of each point is less than half of its value. **a** EDD versus μ_0 , when all $[\mu]_i = \mu_0$; **b** EDD versus p when we randomly select $100p\%$ entries $[\mu]_i$ to be 1 and set the other entries to be 0; the smallest value of p is 0.05

projections in Figs. 5 and 6, this example shows that projection with 0-1 entries can have little performance loss in some cases, and it is still a viable candidate since such projection means a simpler measurement scheme.

5.4 Comparison with multivariate CUSUM

We compare our sketching method with a benchmark adapted from the conventional multivariate CUSUM procedure [42] for the sketches. A main difference is that in multivariate CUSUM, one needs a prescribed post-change mean vector (which is set to be an all-one vector in our example), rather than estimate it as the GLR statistic does. Hence, its performance may be affected by parameter misspecification. We compare the performance again in two

Table 4 A being a fixed Gaussian random matrix

p	0.1	0.2	0.3	0.5	0.7
M_{\min}	300	200	150	100	50

Minimum M_{\min} required for various sparsity setting with parameter p as shown in Fig. 5b. $N = 500$, $w = 200$, and $100p\%$ of entries $[\mu]_i = 1$

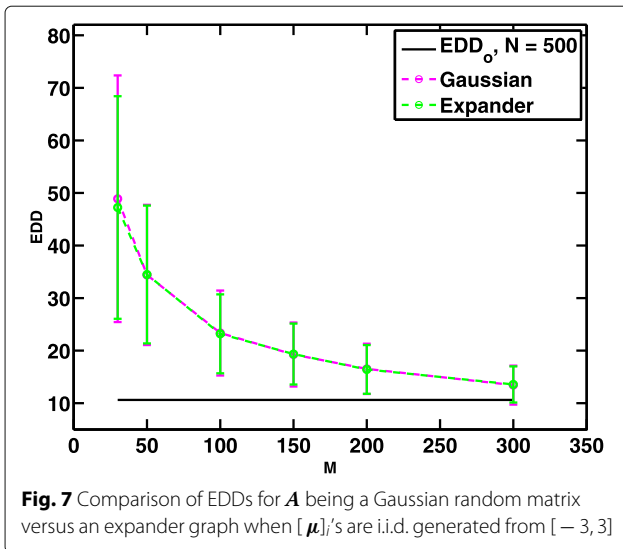


Fig. 7 Comparison of EDDs for \mathbf{A} being a Gaussian random matrix versus an expander graph when $[\mu]_i$'s are i.i.d. generated from $[-3, 3]$

settings, when all $[\mu]_i$ are equal to a constant and when 100p% entries of the post-change mean vector are positive valued. In Fig. 9, the log-GLR-based sketching procedure performs much better than the multivariate CUSUM.

6 Examples for real applications

6.1 Solar flare detection

We use our method to detect a solar flare in a video sequence from the Solar Data Observatory (SDO)¹. Each frame is of size 232×292 pixels, which results in an ambient dimension $N = 67,744$. In this example, the normal frames are slowly drifting background sun surfaces, and the anomaly is a much brighter transient solar flare emerges at $t = 223$. Figure 10a is a snapshot of the original SDO data at $t = 150$ before the solar flare emerges, and Fig. 10b is a snapshot at $t = 223$ when the solar flare emerges as a brighter curve in the middle of the image. We preprocess the data by tracking and removing the slowly changing background with the MOUSSE algorithm [43] to obtain tracking residuals. The Gaussianity for the residuals, which corresponds to our \mathbf{x}_t , is verified by the Kolmogorov-Smirnov test. For instance, the p value is 0.47 for the signal at $t = 150$, which indicates that the Gaussianity is a reasonable assumption.

We apply the sketching procedure with fixed projection to the MOUSSE residuals, choosing the sketching matrix \mathbf{A} to be an M -by- N Gaussian random matrix with entries i.i.d. $\mathcal{N}(0, 1/N)$. Note that the signal is deterministic in this case. To evaluate our method, we run the procedure 500 times, each time using a different random Gaussian matrix as the fixed projection \mathbf{A} . Figure 11 shows the error bars of the EDDs from 500 runs. As M increases, both the means and standard deviations of the EDDs decrease.

¹The video can be found at <http://nislabs.ee.duke.edu/MOUSSE/>. The Solar Object Locator for the original data is SOL2011-04-30T21-45-49L061C108.

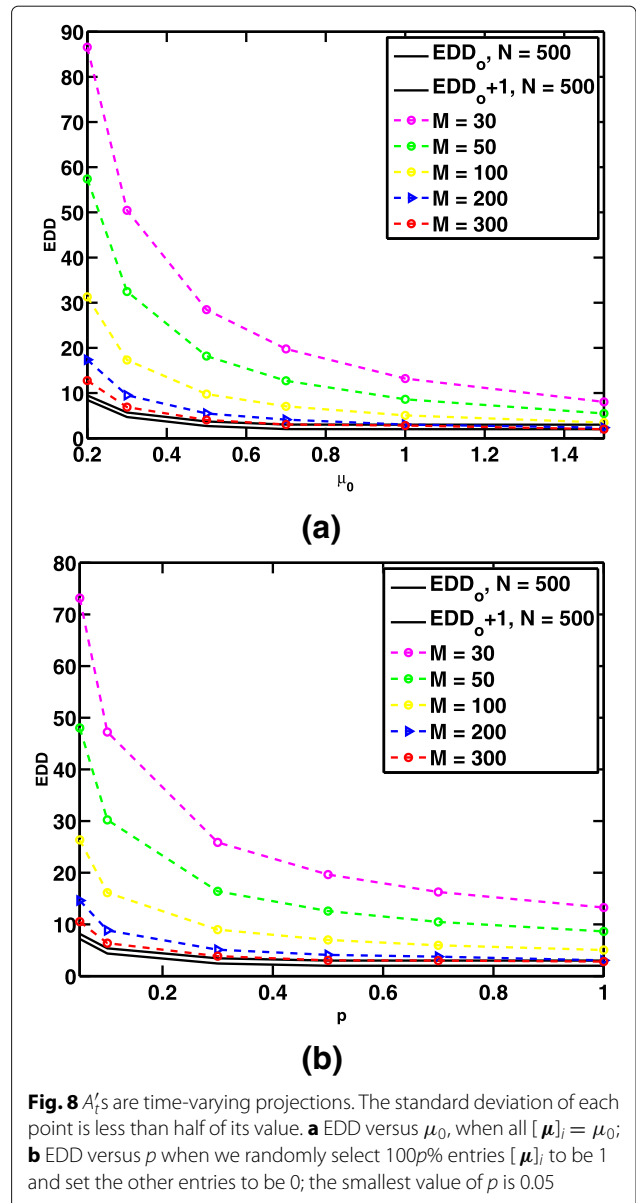
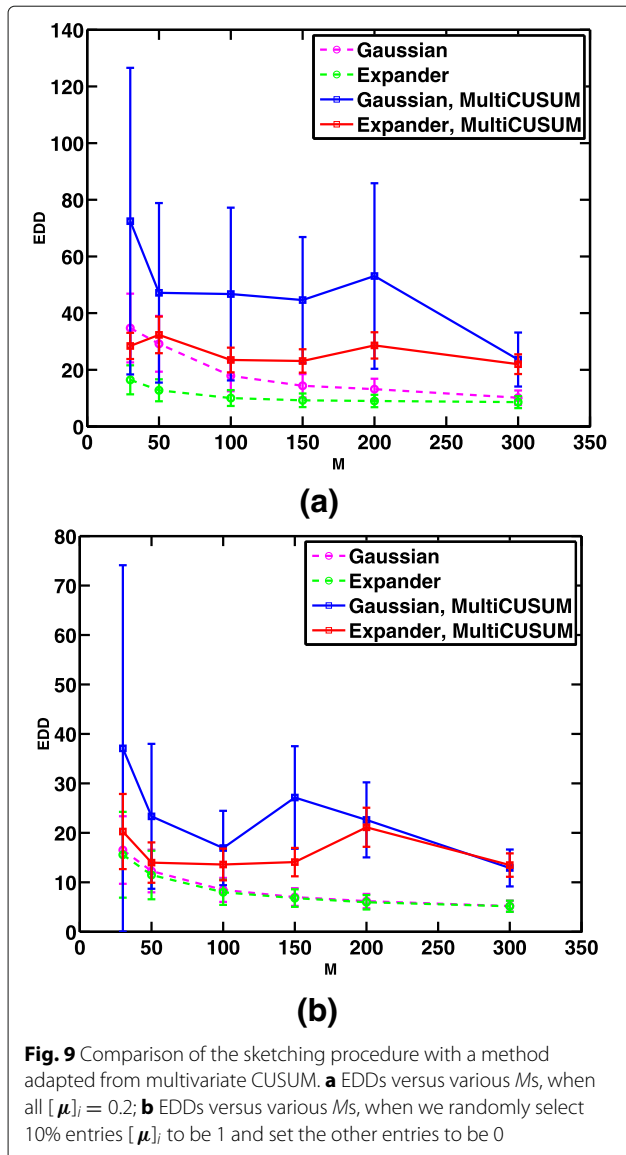


Fig. 8 \mathbf{A}'_t 's are time-varying projections. The standard deviation of each point is less than half of its value. **a** EDD versus μ_0 , when all $[\mu]_i = \mu_0$; **b** EDD versus p when we randomly select 100p% entries $[\mu]_i$ to be 1 and set the other entries to be 0; the smallest value of p is 0.05

When M is larger than 750, EDD is often less than 3, which means that our sketching detection procedure can reliably detect the solar flare with only 750 sketches. This is a significant reduction, and the dimensionality reduction ratio is $750/67,744 \approx 0.01$.

6.2 Change-point detection for power systems

Finally, we present a synthetic example based on the real power network topology. We consider the Western States Power Grid of the USA, which consists of 4941 nodes and 6594 edges. The minimum degree of a node in the network is 1, as shown in Fig. 12. The nodes represent generators, transformers, and substations, and edges represent high-voltage transmission lines between them [44].

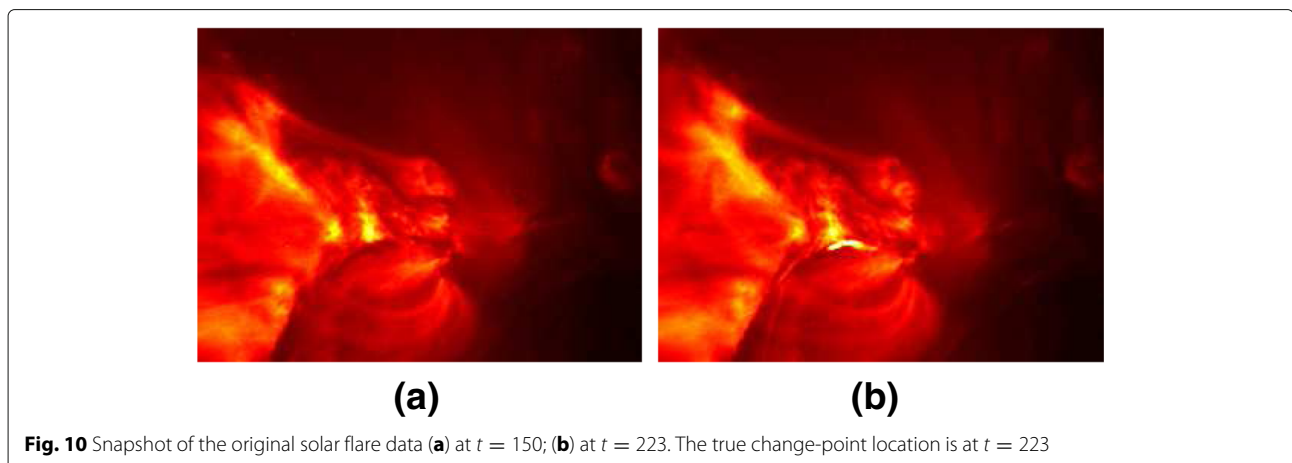


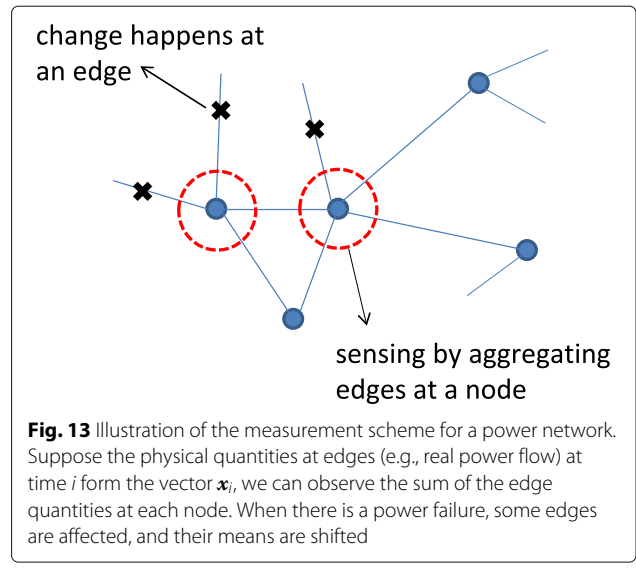
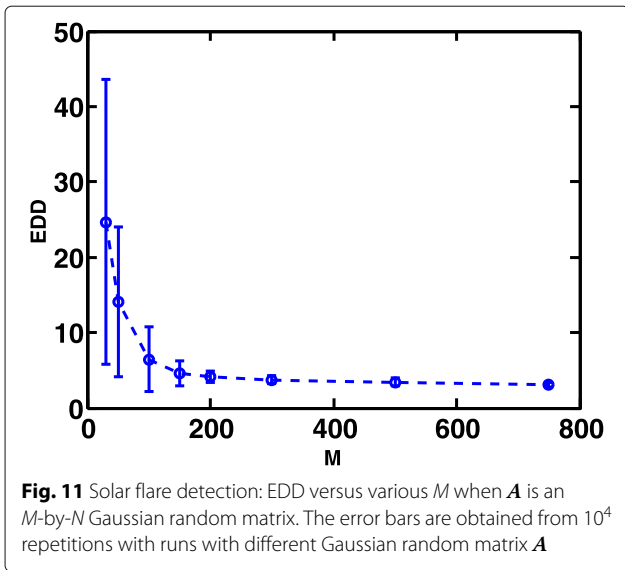
Note that the graph is sparse and that there are many “communities” which correspond to densely connected subnetworks.

In this example, we simulate a situation for power failure over this large network. Assume that at each time, we may observe the real power injection at an edge. When the power system is in a steady state, the observation is the true state plus Gaussian observation noise [45]. We may estimate the true state (e.g., using techniques in [45]), subtract it from the observation vector, and treat the residual vector as our signal x_i , which can be assumed to be i.i.d. standard Gaussian. When a failure happens in a power system, there will be a shift in the mean for a small number of affected edges, since in practice, when there is a power failure, usually only a small part of the network is affected simultaneously.

To perform sketching, at each time, we randomly choose M nodes in the network and measure the sum of the quantities over all attached edges as shown in Fig. 13. This corresponds to A_t^i s with $N = 6594$ and various $M < N$. Note that in this case, our projection matrix is a 0-1 matrix whose structure is constrained by the network topology. Our example is a simplified model for power networks and aims to shed some light on the potential of our method applied to monitoring real power networks.

In the following experiment, we assume that on average, 5% of the edges in the network increase by μ_0 . Set the threshold b such that the ARL is 5000. Figure 14 shows the simulated EDD versus an increasing signal strength μ_0 . Note that the EDD from using a small number of sketches is quite small if μ_0 is sufficiently large. For example, when $\mu_0 = 4$, one may detect the change by observing from only $M = 100$ sketches (when the EDD is increased only by one sample), which is a significant dimensionality reduction with a ratio of $100/4941 \approx 0.02$.



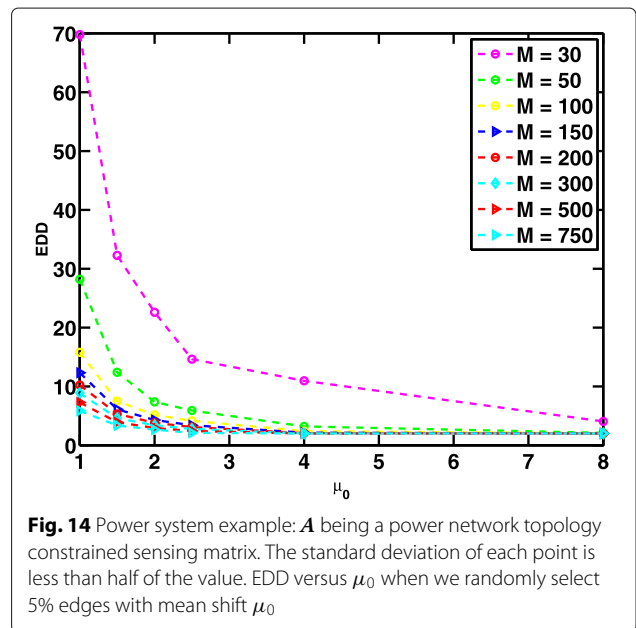
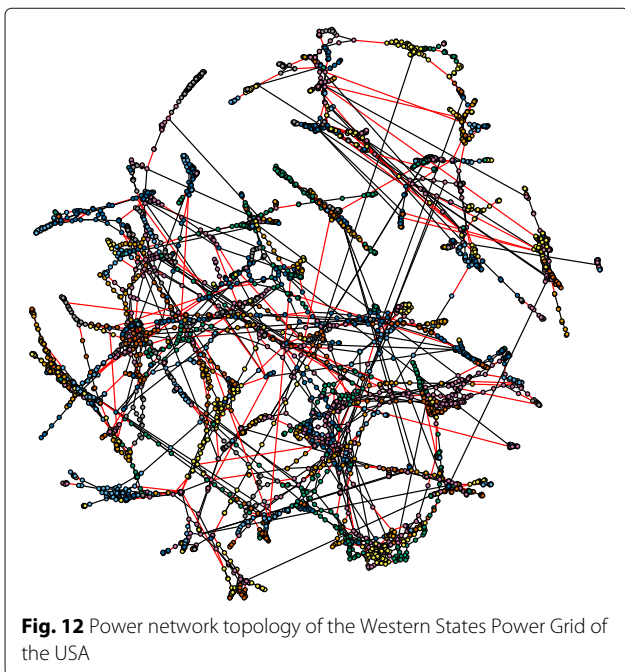


7 Conclusion

In this paper, we studied the problem of sequential change-point detection when the observations are linear projections of the high-dimensional signals. The change-point causes an *unknown* shift in the mean of the signal, and one would like to detect such a change as quickly as possible. We presented new sketching procedures for fixed and time-varying projections, respectively. Sketching is used to reduce the dimensionality of the signal and thus computational complexity; it also reduces data collection and transmission burdens for large systems.

The sketching procedures were derived based on the generalized likelihood ratio statistic. We analyzed the theoretical performance of our procedures by deriving approximations to the average run length (ARL) when there is no change, and the expected detection delay (EDD) when there is a change. Our approximations were shown to be highly accurate numerically and were used to understand the effect of sketching.

We also characterized the relative performance of the sketching procedure compared to that without sketching. We specifically studied the relative performance measure for fixed Gaussian random projections and expander graph projections. Our analysis and numerical examples



showed that the performance loss due to sketching could be quite small in a big regime when the signal strength and the dimension of sketches M are sufficiently large. Our result can also be used to find the minimum required M given a worst-case signal and a target ARL. In other words, we can determine the region where sketching results in little performance loss. We demonstrate the good performance of our procedure using numerical simulations and two real-world examples for solar flare detection and failure detection in power networks.

On a high level, although after sketching, the Kullback-Leibler (K-L) divergence becomes smaller, the *threshold* b for the same ARL also becomes smaller. For instance, for Gaussian matrix, the reduction in K-L divergence is compensated by the reduction of the threshold b for the same ARL, because the factor that they are reduced by are roughly the same. This leads to the somewhat counter-intuitive result that the EDDs with and without sketching turns to be similar in this big regime.

Thus far, we have assumed that the data streams are independent. In practice, if the data streams are dependent on a *known* covariance matrix Σ , we can whiten the data streams by applying a linear transformation $\Sigma^{-1/2}x_t$. Otherwise, the covariance matrix Σ can also be estimated using a training stage via regularized maximum likelihood methods (see [46] for an overview). Alternatively, we may estimate the covariance matrix Σ' of the sketches $A\Sigma A^T$ or $A_t\Sigma A_t^T$ directly, which requires fewer samples to estimate due to the lower dimensionality of the covariance matrix. Then, we can build statistical change-point detection procedures using Σ' (similar to what has been done for the projection Hotelling control chart in [19]), which we leave for future work. Another direction of future work is to accelerate the computation of sketching using techniques such as those in [47].

Appendix 1: Proofs

We start by deriving the ARL and EDD for the sketching procedure.

Proofs for Theorems 1 and 2 This analysis demonstrates that the sketching procedure corresponds to the so-called mixture procedure (cf. T_2 in [15]) in a special case of $p_0 = 1$, M sensors, and the post-change mean vector is $V^T\mu$. In [15], Theorem 1, it was shown that the ARL of the mixture procedure with parameter $p_0 \in [0, 1]$ and M sensors is given by

$$\mathbb{E}^\infty\{T\} \sim H(M, \theta_0) / \underbrace{\int_{[2M\gamma(\theta_0)/m_1]^{1/2}}^{[2M\gamma(\theta_0)/m_0]^{1/2}} \gamma v^2(\gamma) d\gamma}_{c'(M, b, w)}, \quad (36)$$

where the detection statistic will search within a time window $m_0 \leq t - k \leq m_1$. Let $g(x, p_0) = \log(1 - p_0 + p_0 e^{x^2/2})$.

Then, $\psi(\theta) = \log \mathbb{E}\{e^{\theta g(U, p_0)}\}$ is the log moment generating function (MGF) for $g(U, p_0)$, $U \sim \mathcal{N}(0, 1)$, θ_0 is the solution to $\dot{\psi}(\theta) = b/M$,

$$H(M, \theta) = \frac{\theta [2\pi \ddot{\psi}(\theta)]^{1/2}}{\gamma(\theta) M^{1/2}} \exp\{M[\theta \dot{\psi}(\theta) - \psi(\theta)]\}, \quad (37)$$

and

$$\gamma(\theta) = \frac{1}{2} \theta^2 \mathbb{E}\{[\dot{g}(U, p_0)]^2 \exp[\theta g(U, p_0) - \psi(\theta)]\}.$$

Note that U^2 is χ_1^2 distributed, whose MGF is given by $\mathbb{E}\{e^{\theta U^2}\} = 1/\sqrt{1 - 2\theta}$. Hence, when $p_0 = 1$,

$$\psi(\theta) = \log \mathbb{E}\{e^{\theta U^2/2}\} = -\frac{1}{2} \log(1 - \theta).$$

The first-order and second-order derivative of the log MGF are given by, respectively,

$$\dot{\psi}(\theta) = \frac{1}{2(1 - \theta)}, \quad \ddot{\psi}(\theta) = \frac{1}{2(1 - \theta)^2} \quad (38)$$

Set $\dot{\psi}(\theta_0) = b/M$. We obtain the solution that $1 - \theta_0 = M/(2b)$, and $\theta_0 = 1 - M/(2b)$. Hence, $\ddot{\psi}(\theta_0) = 2b^2/M^2$. We have $g(x, 1) = x^2/2$, and $\dot{g}(x, 1) = x$.

$$\begin{aligned} \gamma(\theta) &= \frac{\theta^2}{2} \mathbb{E}\left\{U^2 e^{\frac{\theta U^2}{2}}\right\} e^{\log \sqrt{1 - \theta}} \\ &= \frac{\theta^2}{2} \cdot \frac{1}{(1 - \theta)^{3/2}} \cdot \sqrt{1 - \theta} = \frac{\theta^2}{2(1 - \theta)}, \end{aligned}$$

where

$$\begin{aligned} \mathbb{E}\left\{U^2 e^{\frac{\theta U^2}{2}}\right\} &= \frac{1}{\sqrt{2\pi}} \int x^2 e^{\frac{\theta x^2}{2}} e^{-\frac{x^2}{2}} dx \\ &= \frac{1}{\sqrt{2\pi}} \int x^2 e^{-\frac{x^2}{2(1 - \theta)}} dx = \frac{1}{(1 - \theta)^{3/2}}. \end{aligned}$$

Combining the above, we have that the ARL of the sketching procedure is given by

$$\begin{aligned} \mathbb{E}^\infty\{T\} &= \frac{\theta_0 \left[2\pi \cdot \frac{1}{2(1 - \theta_0)^2}\right]^{1/2}}{c'(M, b, w) \frac{\theta_0^2}{2(1 - \theta_0)} \sqrt{M}} e^{\frac{M\theta_0}{2(1 - \theta_0)}} (1 - \theta_0)^{M/2} + o(1) \\ &= \frac{\sqrt{\pi}}{c'(M, b, w) \theta_0 \sqrt{M}} \frac{2}{e^{\frac{M\theta_0}{2(1 - \theta_0)}}} (1 - \theta_0)^{M/2} + o(1). \end{aligned} \quad (39)$$

Next, using the fact that $1/(1 - \theta_0) = 2b/M$, we have that the two terms in the above expression can be written as

$$\frac{M\theta_0}{2(1 - \theta_0)} = \frac{M\theta_0}{2} \frac{2b}{M} = \theta_0 b, \quad (1 - \theta_0) = \frac{M}{2b},$$

then (39) becomes

$$\begin{aligned} \mathbb{E}^\infty\{T\} &= \frac{\sqrt{\pi}}{c'(M, b, w)} \frac{2}{\theta_0 \sqrt{M}} e^{\theta_0 b} \left(\frac{M}{2b}\right)^{\frac{M}{2}} + o(1) \\ &= \frac{2\sqrt{\pi}}{c'(M, b, w)} \frac{1}{\sqrt{M}} \frac{1}{1 - \frac{M}{2b}} e^{(b - \frac{M}{2})} \left(\frac{M}{2b}\right)^{\frac{M}{2}} + o(1) \\ &= \frac{2\sqrt{\pi}}{c'(M, b, w)} \frac{1}{\sqrt{M}} \frac{1}{1 - \frac{M}{2b}} \left(\frac{M}{2}\right)^{\frac{M}{2}} b^{-\frac{M}{2}} e^{b - \frac{M}{2}} + o(1). \end{aligned}$$

Finally, note that we can also write

$$\gamma(\theta_0) = \theta_0^2 / [2(1 - \theta_0)] = (1 - M/(2b))^2 / (M/b),$$

and the constant is

$$\begin{aligned} c'(M, b, w) &= \int_{[2M\gamma(\theta_0)/w]^{1/2}}^{[2M\gamma(\theta_0)]^{1/2}} yv^2(y)dy \\ &= \int_{\sqrt{\frac{2b}{w}(1 - \frac{M}{2b})}}^{\sqrt{2b(1 - \frac{M}{2b})}} yv^2(y)dy. \end{aligned} \tag{40}$$

We are done deriving the ARL. The EDD can be derived by applying Theorem 2 of [15] in the case where $\Delta = \|\mathbf{V}^\top \boldsymbol{\mu}\|$, the number of sensors is M , and $p_0 = 1$. \square

The following proof is for the Gaussian random matrix \mathbf{A} .

Proof of Theorem 4 It follows from (32), and a standard result concerning the distribution function of the beta distribution ([48], 26.5.3) that

$$\mathbb{P}\{\Gamma \leq b\} = I_b\left(\frac{M}{2}, \frac{N - M}{2}\right), \tag{41}$$

where I is the regularized incomplete beta function (RIBF) ([48], 6.6.2). We first prove the lower bound in (34). Assuming $N \rightarrow \infty$ such that (33) holds, we may combine (41) with ([49], Theorem 4.18) to obtain

$$\begin{aligned} \lim_{(M, N) \rightarrow \infty} \frac{1}{N} \ln \mathbb{P}\{\Gamma \leq \delta - \epsilon\} \\ = - \left[\delta \ln \left(\frac{\delta}{\delta - \epsilon} \right) + (1 - \delta) \ln \left(\frac{1 - \delta}{1 - \delta + \epsilon} \right) \right] = -c < 0, \end{aligned}$$

from which it follows that there exists \tilde{N} such that for all $N \geq \tilde{N}$,

$$\frac{1}{N} \ln \mathbb{P}\{\Gamma \leq \delta - \epsilon\} < -\frac{c'}{2},$$

which rearranges to give

$$\mathbb{P}\{\Gamma \leq \delta - \epsilon\} < e^{-\frac{c'N}{2}},$$

which proves the lower bound in (34). To prove the upper bound, it follows from (41) and a standard property of the RIBF ([48], 6.6.3) that

$$\mathbb{P}\{\Gamma \geq b\} = I_{1-b}\left(\frac{N - M}{2}, \frac{M}{2}\right). \tag{42}$$

Assuming $N \rightarrow \infty$ such that (33) holds, we may combine (42) with ([49], Theorem 4.18) to obtain

$$\begin{aligned} \lim_{(M, N) \rightarrow \infty} \frac{1}{N} \ln \mathbb{P}\{\Gamma \geq \delta + \epsilon\} \\ = - \left[(1 - \delta) \ln \left(\frac{1 - \delta}{1 - \delta - \epsilon} \right) + \delta \ln \left(\frac{\delta}{\delta + \epsilon} \right) \right] = -d < 0, \end{aligned}$$

and the argument now proceeds analogously to that for the lower bound. \square

Lemma 3 *If a 0-1 matrix \mathbf{A} has constant column sum d , for every non-negative vector \mathbf{x} such that $[\mathbf{x}]_i \geq 0$, we have*

$$\|\mathbf{A}\mathbf{x}\|_2 \geq \sqrt{d}\|\mathbf{x}\|_2. \tag{43}$$

Proof of Lemma 3 Below, $A_{ij} = [\mathbf{A}]_{ij}$.

$$\begin{aligned} \|\mathbf{A}\mathbf{x}\|_2^2 &= \sum_{i=1}^M \left(\sum_{j=1}^N A_{ij}x_j \right)^2 \\ &\geq \sum_{i=1}^M \sum_{j=1}^N (A_{ij}x_j)^2 = d\|\mathbf{x}\|_2^2. \end{aligned}$$

\square

Lemma 4 (Bounding $\sigma_{\max}(\mathbf{A})$) *If \mathbf{A} corresponds to a (s, ϵ) -expander with regular degree d and regular left degree c , for any nonnegative vector \mathbf{x} ,*

$$\frac{\|\mathbf{A}\mathbf{x}\|_2}{\|\mathbf{x}\|_2} \leq d\sqrt{\frac{N}{M}}, \tag{44}$$

thus,

$$\sigma_{\max}(\mathbf{A}) \leq d\sqrt{\frac{N}{M}}. \tag{45}$$

Proof of Lemma 4 For any nonnegative vector \mathbf{x} ,

$$\begin{aligned} \|\mathbf{A}\mathbf{x}\|_2^2 &= \sum_{i=1}^M \left(\sum_{j=1}^N A_{ij}x_j \right)^2 \\ &= \sum_{i=1}^M \left(\sum_{j=1}^N (A_{ij}x_j)^2 + \sum_{j=1}^N \sum_{l=1, l \neq j}^N (A_{ij}A_{il}x_jx_l) \right) \\ &\leq \sum_{i=1}^M \left(\sum_{j=1}^N (A_{ij}x_j)^2 + \sum_{j=1}^N \sum_{l=1, l \neq j}^N \frac{A_{ij}A_{il}}{2} (x_j^2 + x_l^2) \right) \end{aligned}$$

$$\begin{aligned}
&= \sum_{i=1}^M \sum_{j=1}^N \sum_{l=1}^N \frac{A_{ij}A_{il}}{2} (x_j^2 + x_l^2) \\
&= \sum_{j=1}^N \sum_{l=1}^N \sum_{i=1}^M \frac{A_{ij}A_{il}}{2} (x_j^2 + x_l^2) \\
&\leq \sum_{j=1}^N dc(x_j)^2 \tag{46}
\end{aligned}$$

$$= \frac{d^2N}{M} \|\mathbf{x}\|_2^2. \tag{47}$$

Above, (46) holds since for a given column j , $A_{ij} = 1$ holds for exactly d rows. And for each row i of these d rows, $A_{il} = 1$ for exactly c columns with $l \in \{1, \dots, p\}$; (47) holds since $dN = Mc$. Finally, from the definition of σ_{\max} , (45) holds. \square

Proof for Theorem 5 Note that

$$\begin{aligned}
\Delta &= (\boldsymbol{\mu}^\top V V^\top \boldsymbol{\mu})^{1/2} = (\mathbf{m} \mathbf{u}^\top \mathbf{A}^\top \mathbf{U} \boldsymbol{\Sigma}^{-2} \mathbf{U}^\top \mathbf{A} \boldsymbol{\mu})^{1/2} \\
&\geq \sigma_{\max}^{-1}(\mathbf{A}) \|\mathbf{U}^\top \mathbf{A} \boldsymbol{\mu}\|_2 = \sigma_{\max}^{-1}(\mathbf{A}) \|\mathbf{A} \boldsymbol{\mu}\|_2, \tag{48}
\end{aligned}$$

where $\sigma_{\max} = \sigma_{\max}(\mathbf{A})$, and (48) holds since \mathbf{U} is a unitary matrix. Thus, in order to bound Δ , we need to characterize σ_{\max} , as well as $\|\mathbf{A} \boldsymbol{\mu}\|_2$ for a s sparse vector $\boldsymbol{\mu}$. Combining (48) with Lemma 3 and 4, we have that for every nonnegative vector $\boldsymbol{\mu}$, $[\boldsymbol{\mu}]_i \geq 0$,

$$\Delta \geq \frac{1}{d} \sqrt{\frac{M}{N}} \sqrt{d(1-\epsilon)} \|\boldsymbol{\mu}\|_2 = \sqrt{\frac{M(1-\epsilon)}{dN}} \|\boldsymbol{\mu}\|_2. \tag{49}$$

Finally, Lemma 2 characterizes the quantity $[M(1-\epsilon)/(dN)]^{1/2}$ in (49) and establishes the existence of such an expander graph. When \mathbf{A} corresponds to an $(\alpha N, \epsilon)$ expander described in Lemma 2, $\Delta \geq \|\beta \boldsymbol{\mu}\|_2$ for all non-negative signals $[\boldsymbol{\mu}]_i \geq 0$ for some constant α and some constant $\beta = (\rho(1-\epsilon)/d)^{1/2}$. Done. \square

Proof for Corollary 1 We define that $x \triangleq M/b$, then Theorem 1 tells us that when M goes to infinity, we have that

$$\begin{aligned}
\mathbb{E}^\infty\{T\} &= \\
&\frac{2\sqrt{\pi}}{c(M, x, w)} \frac{1}{1 - \frac{x}{2}} \frac{1}{\sqrt{M}} \left(\frac{x}{2}\right)^{\frac{M}{2}} \exp\left(\frac{M}{x} - \frac{M}{2}\right) + o(1), \tag{50}
\end{aligned}$$

where

$$c(M, x, w) = \int_{\sqrt{\frac{2M}{xw}}(1-\frac{x}{2})}^{\sqrt{\frac{2M}{x}}(1-\frac{x}{2})} uv^2(u) du, \tag{51}$$

and

$$v(u) \approx \frac{2/u[\Phi(u/2) - 0.5]}{(u/2)\Phi(u/2) + \phi(u/2)}.$$

Define that $\gamma \triangleq \mathbb{E}^\infty\{T\}$. One claim is that when $M > 24.85$ and $\gamma \in [e^5, e^{20}]$, there exists one $x^* \in (0.5, 2)$ such that (50) holds. Next, we prove the claim.

Define the logarithm of the right-hand side of (50) as follows:

$$\begin{aligned}
p(x) &\triangleq \log(2\sqrt{\pi}) - \log(C(M, x, w)) - \log\left(1 - \frac{x}{2}\right) \\
&\quad + \frac{M}{2} \log \frac{x}{2} + \frac{M}{x} - \frac{M}{2} - \frac{1}{2} \log M.
\end{aligned}$$

Since $v(u) \rightarrow 1$ as $u \rightarrow 0$ and $v(u) \rightarrow \frac{2}{u^2}$ as $u \rightarrow \infty$, we know that $\int_0^\infty uv^2(u) du$ exists. From the numerical integration, we know that $\int_0^\infty uv^2(u) du < 1$. Therefore, $-\log(C(M, x, w)) > 0$. Then,

$$p(0.5) > \left(\frac{3}{2} - \frac{1}{2} \log 4\right) M - \frac{1}{2} \log M + \log(2\sqrt{\pi}) - \log \frac{3}{4}.$$

When $M > 24.85$, we have that $p(0.5) > 20$. Then, when $\gamma < e^{20}$, we have that $p(0.5) - \log \gamma > 0$.

Next, we prove that we can find some $x_0 \in (0.5, 2)$ such that $p(x_0) - \log \gamma < 0$ provided that $\gamma > e^5$. Since $\phi\left(\frac{u}{2}\right) < 0.5$ and

$$0.5 + \frac{1}{\sqrt{2\pi}} \exp\left[-\frac{1}{2} \left(\frac{u}{2}\right)^2\right] \left(\frac{u}{2}\right) \leq \Phi\left(\frac{u}{2}\right) \leq 1,$$

for any $u > 0$. We have that

$$v(u) > \sqrt{\frac{2}{\pi}} \cdot \frac{\exp\left(-\frac{u^2}{8}\right)}{u+1}.$$

Then, we have that for any $u > 0$,

$$uv^2(u) > \frac{2}{\pi} \cdot \frac{u \cdot \exp(-u^2/4)}{(u+1)^2}.$$

We define that x_0 is the solution to the following equation:

$$\sqrt{\frac{2M}{x}} \left(1 - \frac{x}{2}\right) = 1. \tag{52}$$

Then, we have that

$$\begin{aligned}
C(M, x_0, w) &> \frac{2}{\pi} \cdot \int_{1/\sqrt{w}}^1 \frac{u \cdot \exp(-u^2/4)}{(u+1)^2} du \\
&> \frac{2}{\pi} \cdot \int_{1/\sqrt{w}}^1 \frac{u \cdot \exp(-u^2/4)}{4} du \\
&= \frac{1}{\pi} \cdot \left(\exp\left(-\frac{1}{4w}\right) - \exp\left(-\frac{1}{4}\right)\right) \\
&> \frac{1}{\pi} \exp\left(-\frac{1}{4}\right) \cdot \left(\frac{1}{4} - \frac{1}{4w}\right),
\end{aligned}$$

where the second inequality is due to the fact that the upper bound for the integral interval is 1 and the third inequality is due to the fact that $\exp(-x)$ is a convex function. Therefore, we have that

$$-\log C(M, x_0, w) < \log \pi + \frac{1}{4} - \log\left(\frac{1}{4} - \frac{1}{4w}\right)$$

Note that the upper bound above for $-\log C(M, x_0, w)$ is not dependent on M , which is because we choose a x_0 that depends on M . Solving the Eq. (52), we have that

$$x_0 = 2 + \frac{1}{M} - \sqrt{\frac{1}{M^2} + \frac{4}{M}} < 2,$$

and $x_0 \rightarrow 2$ as $M \rightarrow \infty$. By Taylor's expansion, we have that $x_0 = 2 - 2M^{-1/2} + M^{-1} + o(M^{-1})$, or $x_0 = 2 - 2M^{-1/2} + o(M^{-1/2})$. Then, we have that

$$-\log\left(1 - \frac{x_0}{2}\right) = -\log(M^{-1/2}) + o(1),$$

and

$$\begin{aligned} \frac{M}{2} \log \frac{x_0}{2} &= \frac{M}{2} \log(1 - M^{-1/2}) \\ &= \frac{M}{2} \cdot \left(-M^{-1/2} - \frac{1}{2}M^{-1} + o(M^{-1})\right) \\ &= -\frac{1}{2}M^{1/2} - \frac{1}{4} + o(1), \end{aligned}$$

and

$$\begin{aligned} \frac{M}{x_0} &= \frac{M}{2} \cdot \frac{1}{1 - (M^{-1/2} + M^{-1}/2 + o(M^{-1}))} \\ &= \frac{M}{2} \cdot (M^{-1/2} + M^{-1}/2 + o(M^{-1}) \\ &\quad + (M^{-1/2} + M^{-1}/2 + o(M^{-1}))^2 + o(M^{-1})) \\ &= \frac{1}{2}M^{1/2} + \frac{1}{2} + o(1) \end{aligned}$$

Combining the above results, we have that

$$p(x_0) < \log(2\sqrt{\pi}) + \log \pi - \log\left(\frac{1}{4} - \frac{1}{4w}\right) + \frac{1}{2} + o(1). \tag{53}$$

One important observation is that the right-hand side of (53) converges as $M \rightarrow \infty$. In fact, $p(x_0)$ as a function of M is decreasing and converges as $M \rightarrow \infty$. Since we set $w \geq 100$, then for any $M > 24.85$, $p(x_0) < 5$. Therefore, for any $\gamma > e^5$ and any $M > 24.85$, we can find a x_0 close to 2 such that $p(x_0) - \log \gamma < 0$.

Since $p(x)$ is a continuous function, there exists a solution $x^* \in (0.5, 2)$ such that Eq. (50) holds. \square

Proof of Theorem 3 The proof uses a similar argument as that in [15].

By law of large number, when $t - k$ tends to infinity, the following sum converges in probability

$$\frac{1}{t - k} \sum_{i=k+1}^t \mathbb{I}_{in} \xrightarrow{p} r. \tag{54}$$

Moreover, from central limit theorem,

$$\frac{1}{\sqrt{t - k}} \sum_{i=k+1}^t [\mathbf{x}_i]_n (\mathbb{I}_{in} - r) \xrightarrow{d} \mathcal{N}(0, r(1 - r)). \tag{55}$$

So by continuous mapping theorem,

$$\left(\frac{1}{\sqrt{(t - k)r(1 - r)}} \sum_{i=k+1}^t [\mathbf{x}_i]_n (\mathbb{I}_{in} - r)\right)^2 \xrightarrow{d} \chi_1^2, \tag{56}$$

i.e., the squared and scaled version of the sum is asymptotically a χ_1^2 random variable with one degree of freedom. By Slutsky's theorem, combining (54) and (56),

$$\frac{1}{1 - r} \frac{[\sum_{i=k+1}^t \mathbf{x}_i]_n (\mathbb{I}_{in} - r)^2}{\sum_{i=k+1}^t \mathbb{I}_{in}} \xrightarrow{d} \chi_1^2$$

Using Lemma 1 in [50], for $X \sim \chi_1^2$,

$$\begin{aligned} \mathbb{P}\{X \geq 1 + 2\sqrt{\epsilon} + 2\epsilon\} &\leq e^{-\epsilon} \\ \mathbb{P}\{X \leq 1 - 2\sqrt{\epsilon}\} &\leq e^{-\epsilon} \end{aligned}$$

Therefore, with probability at least $1 - 2e^{-\epsilon}$, the difference is bounded by a constant

$$\left(\frac{\sum_{i=k+1}^t [\mathbf{x}_i]_n \mathbb{I}_{in}}{\sqrt{\sum_{i=k+1}^t \mathbb{I}_{in}}} - r \frac{\sum_{i=k+1}^t [\mathbf{x}_i]_n}{\sqrt{\sum_{i=k+1}^t \mathbb{I}_{in}}}\right)^2 < (1 + 2\sqrt{\epsilon} + 2\epsilon)(1 - r).$$

On the hand, by central limit theorem, when $t - k$ tends to infinity,

$$\frac{1}{\sqrt{t - k}} \sum_{i=k+1}^t [\mathbf{x}_i]_n \xrightarrow{d} \mathcal{N}(0, 1).$$

and by law of large number and continuous mapping theorem

$$\left(\frac{\sum_{i=k+1}^t \mathbb{I}_{in}}{t - k}\right)^{-1/2} - \frac{1}{\sqrt{r}} \xrightarrow{p} 0$$

Hence, invoking Slutsky's theorem again, we have

$$\left(\frac{\sum_{i=k+1}^t [\mathbf{x}_i]_n}{\sqrt{\sum_{i=k+1}^t \mathbb{I}_{in}}} - \frac{\sum_{i=k+1}^t [\mathbf{x}_i]_n}{\sqrt{r(t - k)}}\right)^2 \xrightarrow{d} 0$$

Hence, combining the above, by a triangle inequality type of argument, we may conclude that, with high probability, the difference is bounded by a constant c

$$\left(\frac{\sum_{i=k+1}^t [\mathbf{x}_i]_n \mathbb{I}_{in}}{\sqrt{\sum_{i=k+1}^t \mathbb{I}_{in}}} - \sqrt{r} \frac{\sum_{i=k+1}^t [\mathbf{x}_i]_n}{\sqrt{(t - k)}}\right)^2 < c.$$

Hence, to control the ARL for the procedure defined in (21)

$$T_{\{0,1\}} = \inf\left\{t : \max_{t-w \leq k < t} \frac{1}{2} \sum_{n=1}^N \frac{(\sum_{i=k+1}^t [\mathbf{x}_i]_n \mathbb{I}_{in})^2}{\sum_{i=k+1}^t \mathbb{I}_{in}} > b\right\}, \quad (57)$$

one can approximately consider another procedure

$$\tilde{T}_{\{0,1\}} = \inf\left\{t : \max_{t-w \leq k < t} \frac{1}{2} \sum_{n=1}^N U_{n,k,t}^2 > \frac{b}{r}\right\},$$

with

$$U_{n,k,t} \triangleq \frac{\sum_{i=k+1}^t [\mathbf{x}_i]_n}{\sqrt{t-k}},$$

This corresponds to the special case of the mixture procedure with N sensors and all being affected by the change ($p_0 = 1$), except that the threshold is scaled by $1/r$. Hence, we can use the ARL approximation for mixture procedure, which leads to (26). \square

Appendix 2: Justification for EDD of (27)

Proof Below, let $T = T'_{\{0,1\}}$ for simplicity. Define $S_{n,t} = \sum_{i=1}^t [\mathbf{x}_i]_n \xi_{ni}$ for any n and t . To obtain an EDD approximation to $T'_{\{0,1\}}$, first we note that

$$\begin{aligned} \frac{1}{2} \sum_{n=1}^N Z_{n,k,T}^2 &= \frac{1}{2} \sum_{n=1}^N \frac{(\sum_{i=k+1}^T [\mathbf{x}_i]_n \xi_{ni})^2}{r(T-k)} \\ &= \frac{1}{2} \sum_{n=1}^N \frac{(S_{n,T} - S_{n,k})^2}{r(T-k)} \\ &= \sum_{n=1}^N \mu_n [S_{n,T} - S_{n,k} - (T-k)r\mu_n/2] \\ &\quad + \sum_{n=1}^N [S_{n,T} - S_{n,k} - (T-k)r\mu_n]^2 / (2r(T-k)). \end{aligned} \quad (58)$$

Then, we can leverage a similar proof as that to Theorem 2 in [15] to obtain that as $b \rightarrow \infty$,

$$\begin{aligned} &\mathbb{E}^0 \left\{ \max_{0 \leq k < T} \frac{1}{2} \sum_{n=1}^N Z_{n,k,T}^2 \right\} \\ &= \mathbb{E}^0 \left\{ \sum_{n=1}^N \mu_n (S_{n,T} - rT\mu_n/2) \right\} \\ &\quad + \mathbb{E}^0 \left\{ \sum_{n=1}^N (S_{n,T} - Tr\mu_n)^2 / (2Tr) \right\} \\ &\quad + \mathbb{E}^0 \left\{ \min_{0 \leq k < b^{1/2}} \left(\sum_{n=1}^N \mu_n (S_{n,k} - kr\mu_n/2) \right) \right\} + o(1). \end{aligned} \quad (59)$$

The first term on the right-hand side of (59) is equal to $\mathbb{E}^0 \left\{ T'_{\{0,1\}} \right\} \cdot r \sum_{n=1}^N \mu_n^2 / 2$. Using the fact that random variables $([\mathbf{x}_i]_n \xi_{ni} - r\mu_n) / \sqrt{r}$ are i.i.d. with mean zero and unit variance, together with the Anscombe-Doebelin Lemma [36], we have that as $b \rightarrow \infty$, the second term on the right-hand side of (59) is equal to $N/2 + o(1)$. The third term can be shown to be small similar to [15]. Finally, ignoring the overshoot of the detection statistic exceeding the detection threshold, we can replace the left-hand side of (59) with b . Solving the equation, we obtain the first order approximation of the EDD is given by (27). \square

Abbreviations

ARL: Average run length; CUSUM: Cumulative sum; EDD: Expected detection delay; EWMA: Exponential weighting moving average; GLR: Generalized likelihood ratio; K-L: Kullback-Leibler; PCA: Principal component analysis; SPC: Statistical control charts; SVD: Singular value decomposition

Acknowledgements

We want to thank the anonymous reviewers for their excellent comments, which help greatly improve the paper. We also want to thank Minghe Zhang for the help with some revisions.

Authors' contributions

We present sequential change-point detection procedures based on linear sketches of high-dimensional signal vectors using generalized likelihood ratio (GLR) statistics. Rigorous theoretical analysis and numerical examples on simulated and real data are also presented. YX came up with original formulation and analysis. YC performed numerical examples and additional theoretical analysis. AT and MW also contributed to theoretical analysis. All authors read and approved the final manuscript.

Funding

This work is partially supported by NSF grants CCF-1442635 and CMMI-1538746 and an NSF CAREER Award CCF-1650913.

Availability of data and materials

The data for solar flare image is at http://www2.isye.gatech.edu/~yxie77/data_193.mat. The topology of the power network can be downloaded at <http://networkdata.ics.uci.edu/data/power/> [44].

Competing interests

The authors declare that they have no competing interests.

Author details

¹H. Milton Stewart School of Industrial and Systems Engineering, Georgia Institute of Technology, Atlanta, GA, USA. ²National Physical Laboratory, Teddington, UK. ³Department of Electrical, Computer & Systems Engineering, Rensselaer Polytechnic Institute, Troy, NY, USA.

Received: 21 January 2019 Accepted: 21 July 2019

Published online: 09 September 2019

References

1. Y. Xie, M. Wang, A. Thompson, in *Global Conference on Signal and Information Processing (GlobalSIP)*. Sketching for sequential change-point detection, (Orlando, 2015), pp. 78–82
2. A. Tartakovsky, I. Nikiforov, M. Basseville, *Sequential analysis: Hypothesis Testing and Changepoint Detection*. (Chapman and Hall/CRC, 2014)
3. HV. Poor, O. Hadjililadis, *Quickest detection*. (Cambridge University Press, Cambridge, 2008)
4. D. P. Woodruff, Sketching as a tool for numerical linear algebra. *Found. Trends. Theor. Comput. Sci.* **10**, 1–157 (2014)
5. G. Dasarthy, P. Shah, B. N. Bhaskar, R. Nowak, in *Communication, Control, and Computing (Allerton)*, 2012 50th Annual Allerton Conference on. Covariance sketching (IEEE, Monticello, 2012), pp. 1026–1033

6. Y. Chi, in *Signal Processing Conference (EUSIPCO), 2016 24th European*. Kronecker covariance sketching for spatial-temporal data (IEEE, Budapest, 2016), pp. 316–320
7. Y. Wang, H.-Y. Tung, A. J. Smola, A. Anandkumar, in *Advances in Neural Information Processing Systems*. Fast and guaranteed tensor decomposition via sketching, (Montreal, 2015), pp. 991–999
8. A. Alaoui, M. W. Mahoney, in *Advances in Neural Information Processing Systems*. Fast randomized kernel ridge regression with statistical guarantees (Montreal, 2015), pp. 775–783
9. Y. Bachrach, R. Herbrich, E. Porat, in *International Symposium on String Processing and Information Retrieval*. Sketching algorithms for approximating rank correlations in collaborative filtering systems (Springer, New York, 2009), pp. 344–352
10. G. Raskutti, M. Mahoney, in *International Conference on Machine Learning*. Statistical and algorithmic perspectives on randomized sketching for ordinary least-squares, (Lille, 2015), pp. 617–625
11. P. Indyk, in *Proceedings of the nineteenth annual ACM-SIAM symposium on Discrete algorithms Society for Industrial and Applied Mathematics*. Explicit constructions for compressed sensing of sparse signals, (San Francisco, 2008), pp. 30–33
12. D. Siegmund, B. Yakir, N. R. Zhang, Detecting simultaneous variant intervals in aligned sequences. *Ann. Appl. Stat.* **5**(2A), 645–668 (2011)
13. G. K. Atia, Change detection with compressive measurements. *Sig. Process. Lett. IEEE.* **22**(2), 182–186 (2015)
14. D. Siegmund, E. S. Venkatraman, Using the generalized likelihood ratio statistic for sequential detection of a change-point. *Ann. Stat.* **23**(1), 255–271 (1995)
15. Y. Xie, D. Siegmund, Sequential multi-sensor change-point detection. *Ann. Stat.* **41**(2), 670–692 (2013)
16. G. C. Runger, Projections and the U-squared multivariate control chart. *J. Qual. Technol.* **28**(3), 313–319 (1996)
17. O. Bodnar, W. Schmid, Multivariate control charts based on a projection approach. *Allg. Stat. Arch.* **89**, 75–93 (2005)
18. E. Skubalska-Rafajlowicz, Random projections and Hotelling's T2 statistics for change detection in high-dimensional data analysis. *Int. J. Appl. Math. Comput. Sci.* **23**(2), 447–461 (2013)
19. E. Skubalska-Rafajlowicz, in *Stochastic models, statistics and their applications*. Change-point detection of the mean vector with fewer observations than the dimension using instantaneous normal random projections, vol. 122 (Springer Proc. Math. Stat, New York, 2015)
20. D. C. Montgomery, *Introduction to statistical quality control*. (Wiley, 2008)
21. M. A. Davenport, P. T. Boufounos, M. B. Wakin, R. G. Baraniuk, Signal processing with compressive measurements. *Sel. Top. Sig. Process. IEEE. J.* **4**(2), 445–460 (2010)
22. E. Arias-Castro, et al., Detecting a vector based on linear measurements. *Electron. J. Stat.* **6**, 547–558 (2012)
23. J. Geng, W. Xu, L. Lai, in *Information Theory (ISIT), 2013 IEEE International Symposium on, Istanbul, Turkey*. Quickest search over multiple sequences with mixed observations, (2013), pp. 2582–2586
24. W. Xu, L. Lai, in *Communication, Control, and Computing (Allerton), 2013 51st Annual Allerton Conference on*. Compressed hypothesis testing: to mix or not to mix? (IEEE, Monticello, 2013), pp. 1443–1449
25. Z. Harchaoui, E. Moulines, F. R. Bach, in *Advances in Neural Information Processing Systems*. Kernel change-point analysis, (Vancouver, 2009), pp. 609–616
26. S. Arlot, A. Celisse, Z. Harchaoui, Kernel change-point detection (2012). arXiv preprint arXiv:1202.3878
27. Y. C. Chen, T. Banerjee, A. D. Domínguez-García, V. V. Veeravalli, Quickest line outage detection and identification. *Power. Syst. IEEE. Trans.* **31**(1), 749–758 (2016)
28. D. Mishin, K. Brantner-Magee, F. Czako, A. S. Szalay, in *High Performance Extreme Computing Conference (HPEC), 2014 IEEE*. Real time change point detection by incremental PCA in large scale sensor data (IEEE, Waltham, 2014), pp. 1–6
29. F. Tsung, K. Wang, in *Frontier in Statistical Quality Control*, pp. 19–35. Adaptive charting techniques: literature review and extensions (Springer-Verlag, New York, 2010)
30. J. Chen, S.-H. Kim, Y. Xie, S3T: an efficient score-statistic for spatio-temporal surveillance (2017). arXiv:1706.05331
31. W. Xu, B. Hassibi, in *Info. Theory Workshop*. Efficient compressive sensing with deterministic guarantees using expander graphs, (2007)
32. Y. Chen, C. Suh, A. J. Goldsmith, in *Information Theory (ISIT), 2015 IEEE International Symposium on, Hong Kong*. Information recovery from pairwise measurements: a Shannon-theoretic approach (IEEE, Hong Kong, 2015), pp. 2336–2340
33. A. K. Massimino, M. A. Davenport, in *Proc. Workshop on Signal Processing with Adaptive Sparse Structured Representations (SPARS)*. One-bit matrix completion for pairwise comparison matrices, (Lausanne, 2013)
34. J. E. Jackson, *A user's guide to principle components*. (Wiley, New York, 1991)
35. L. Balzano, S. J. Wright, Local convergence of an algorithm for subspace identification from partial data. *Found. Comput. Math.* **15**(5), 1279–1314 (2015)
36. D. Siegmund, *Sequential Analysis: Test and Confidence Intervals*. (Springer, New York, 1985)
37. D. Siegmund, B. Yakir, *The statistics of gene mapping*. (Springer, New York, 2007)
38. H. Ruben, The volume of an isotropic random parallelotope. *J. Appl. Probab.* **16**(1), 84–94 (1979)
39. P. Frankl, H. Maehara, Some geometric applications of the beta distribution. *Ann. Inst. Statist. Math.* **42**(3), 463–474 (1990)
40. D. Burshtein, G. Miller, Expander graph arguments for message-passing algorithms. *IEEE Trans. Inf. Theory.* **47**(2), 782–790 (2001)
41. E. J. Candes, The restricted isometry property and its implications for compressed sensing. *Compte Rendus de l'Academie des Sciences, Paris, Serie I.* **342**, 589–592 (2008)
42. W. H. Woodall, M. M. Ncube, Multivariate CUSUM quality-control procedures. *Technometrics.* **27**(3), 285–292 (1985)
43. Y. Xie, J. Huang, R. Willett, Change-point detection for high-dimensional time series with missing data. *Sel. Top. Sig. Process. IEEE. J.* **7**(1), 12–27 (2013)
44. D. J. Watts, S. H. Strogatz, Collective dynamics of 'small-world' networks. *Nature.* **393**(6684), 440–442 (1998)
45. A. Abur, A. G. Exposito, *Power system state estimation: Theory and Implementation*. (CRC Press, 2004)
46. J. Fan, Y. Liao, H. Liu, An overview on the estimation of large covariance and precision matrices. *Econ. J.* **19**(1), C1–C32 (2016)
47. M. Kapralov, V. K. Potluru, D. P. Woodruff, in *International Conference on Machine Learning*. How to fake multiply by a Gaussian matrix, (New York, 2016)
48. M. Abramowitz, I. Stegun, *A handbook of mathematical functions, with formulas, graphs and mathematical tables*, 10th. (Dover, New York, 1964)
49. A. Thompson, Quantitative analysis of algorithms for compressed signal recovery. PhD thesis, School of Mathematics, University of Edinburgh (2012)
50. B. Laurent, P. Massart, Adaptive estimation of a quadratic functional by model selection. *Ann. Stat.* **28**(5), 1302–1338 (2000)

Publisher's Note

Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.

Submit your manuscript to a SpringerOpen[®] journal and benefit from:

- Convenient online submission
- Rigorous peer review
- Open access: articles freely available online
- High visibility within the field
- Retaining the copyright to your article

Submit your next manuscript at ► [springeropen.com](https://www.springeropen.com)