



HAL
open science

Sequential detection of a temporary change in multivariate time series

Victor Watson, François Septier, Patrick Armand, Christophe Duchenne

► **To cite this version:**

Victor Watson, François Septier, Patrick Armand, Christophe Duchenne. Sequential detection of a temporary change in multivariate time series. *Digital Signal Processing*, 2022, 127, pp.103545. 10.1016/j.dsp.2022.103545 . hal-03727692

HAL Id: hal-03727692

<https://hal.science/hal-03727692v1>

Submitted on 22 Jul 2024

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.



Distributed under a Creative Commons Attribution - NonCommercial 4.0 International License

Sequential detection of a temporary change in multivariate time series

Victor Watson^{a,b}, François Septier^b, Patrick Armand^a,
Christophe Duchenne^a

^aCEA, DAM, DIF, , Arpajon, F-91297, France

^bUniv Bretagne Sud, CNRS UMR 6205, LMBA, Rue André
Lwoff, Vannes, F-56000, France

Abstract

In this work, we aim to provide a new and efficient recursive detection method for temporarily monitored signals. Motivated by the case of the propagation of an event over a field of sensors, we assumed that the change in the statistical properties in the monitored signals can only be temporary. Unfortunately, to our best knowledge, existing recursive and simple detection techniques such as the ones based on the cumulative sum (CUSUM) do not consider the temporary aspect of the change in a multivariate time series. In this paper, we propose a novel simple and efficient sequential detection algorithm, named Temporary-Event-CUSUM (TE-CUSUM). By combining with a new adaptive way to aggregate local CUSUM variables from each data stream, we empirically show that the TE-CUSUM has a very good detection rate in the case of an event passing through a field of sensors in a very noisy environment.

Keywords:

sequential detection, multivariate time series, CUSUM, temporary event

1. Introduction

The multiplication of industrial sites near populated areas increases the danger for populations in cases of an unexpected release of a hazardous compound. Densely populated areas can also be a target for ill-intentioned people who would release some toxic material and cause many victims. In such cases early detection can be crucial. If these sensitive areas are monitored, waiting for the level of toxic compound to be sufficient so that it is

Preprint submitted to Elsevier

March 11, 2022

unambiguously monitored by the sensors can have the consequence of being too late in one's response to a threat. Sequential change-point detection uses the statistics of a data stream to detect an abnormality in a signal while the signal is still low. This means that one could detect a small concentration of a chemical in the air before a level threatening human life is reached. One would detect an abnormal radioactivity level due to an irradiating source, for instance contained in a dirty bomb before explosion. These sequential detection techniques can be used to detect the presence of a pollutant in the air such as in [1]. Early detection allows estimation techniques such as in [2] to start monitoring the data at the right moment and facilitates the convergence to a solution while decreasing the computational cost. It is also used for early seismic detection [3], or early detection of infected people during a pandemic [4], but can also be applied to many other fields such as [5] and [6].

The CUSUM (CUMulative SUM) technique [7] is a powerful univariate sequential change-point detection tool on which are based many detection techniques. The extension of the CUSUM to multivariate cases is not straightforward and has been the object of many considerations in the process-control community (see [8, 9, 10, 11, 12, 13, 14]). Moreover, the issue of the possible non-synchronicity of the monitoring between the sensors in multivariate cases for the detection of temporary change remains a research open question. Indeed, the existing ways to deal with these problems require to lose the recursive computation of the test statistics necessary to trigger detection. The temporary change is not a common consideration in the process-control community as when a process gets out-of-control, it rarely gets back in-control. When we extend the sequential detection technique to some other physical problems such as the ones cited earlier, this back in-control scenario is what we expect as a sensor can be exposed only for a limited duration.

In this paper we propose a new multivariate CUSUM-based technique to deal with temporary changes without losing the recursive computation. Indeed, the Temporary-Event-CUSUM (TE-CUSUM) does not require the change to be permanent or synchronous between data streams to be detected. We believe this novel TE-CUSUM will be relevant to many application cases such as when a pollutant released into the air is dispersed through a wide area and low concentrations are recorded by different sensors at different times. To our best knowledge existing methods would require simultaneous sensors' exposures or would be too computationally expensive methods for their online practical use. We also developed a new adaptive method to combine the local

test statistics so it increases the performances of the TE-CUSUM when the subset of sensor which is affected with the signal is unknown.

This paper is organised as follows. In the second section, we introduce our detection problem and the CUSUM technique as a Generalised-likelihood ratio test (GLRT). We set the principle of the proposed TE-CUSUM and show its equivalence with the CUSUM for univariate cases. In the third section, we extend the model problem to multivariate temporary events and make a quick review of the existing methods for multivariate sequential change-point detection. We then introduce a new strategy that combines local statistics with a novel adaptive censoring method and also demonstrate the efficiency of the TE-CUSUM in multivariate sequential detection cases. The fourth section is a validation test in which we compare the efficiency of the different techniques for monitoring the dispersion of a pollutant over a field of sensors.

2. Sequential change-point detection in univariate time series

The change-point detection problem in univariate time series can be formulated as the following hypothesis test:

$$\begin{aligned} \mathcal{H}_0 : p(x_1, \dots, x_n) &= \prod_{k=1}^n f_0(x_k | \theta_0) \\ \mathcal{H}_1 : p(x_1, \dots, x_n; \nu) &= \prod_{k=1}^{\nu} f_0(x_k | \theta_0) \prod_{k=\nu+1}^n f_1(x_k | \theta_1) \end{aligned} \quad (1)$$

This represents a two case scenario, the first one marked by the hypothesis \mathcal{H}_0 for which every sample x_k with $k \leq n$ follows $f_0(x_k | \theta_0)$, the second one stating that there is a time $\nu \in [1; n]$ such that x_k with ($k > \nu$) starts to follow $f_1(x_k | \theta_1)$.

This leads us to the associate likelihood ratio:

$$\Lambda_n^\nu = \prod_{k=\nu+1}^n \frac{f_1(x_k | \theta_1)}{f_0(x_k | \theta_0)} \quad (2)$$

Comparing Λ_n^ν to a threshold allows us to define a statistical test sequentially computed to decide between the two hypothesis.

2.1. Generalised likelihood ratio test and CUSUM

One problem about the likelihood ratio test of Equation (2) is that the knowledge of the change-point ν is needed. In such a situation (unknown parameter in the likelihood distribution), it is common to use a generalised likelihood ratio test (GLRT) [15] which is defined in our problem as:

$$V_n = \max_{0 \leq \nu < n} (\Lambda_n^\nu) = \max_{0 \leq \nu < n} \prod_{k=\nu+1}^n L_k \quad (3)$$

with,

$$L_k = \frac{f_1(x_k|\theta_1)}{f_0(x_k|\theta_0)} \quad (4)$$

while the change point can be estimated with:

$$\hat{\nu} = \operatorname{argmax}_{0 \leq \nu < n} \prod_{k=\nu+1}^n L_k \quad (5)$$

The criterion V_n increases when $L_k > 1$ and decreases when $L_k < 1$. If $k > \nu$, the hypothesis \mathcal{H}_1 is true and the ratio L_k has a better chance to be greater than 1. V_n will overall increase even if monotony is far from guaranteed. We can then compare V_n to a threshold to trigger detection sequentially once a novel observation is received. In such a sequential setting, since we are interested in the quickest detection method, it is also important to consider the detection delay time instead of just the probability of detection [11].

The generalised likelihood ratio in Equation (3) can be rewritten in the following recursive form which allows its integration in online systems:

$$V_n = \max(1, V_{n-1})L_n, \quad n \geq 1, V_0 = 1 \quad (6)$$

2.1.1. CUSUM principle

The CUSUM technique was first introduced by Page in 1954 [7]. This algorithm has been proposed in order to optimise both the detection delay and the average run-length to false alarm (ARL2FA) which is the average time between two false alarms [11]. To compute the ARL2FA, we used the cyclical steady-state ARL described by [16]. This procedure can be seen as a sequential algorithm to recursively compute the GLRT defined by Equation (3). By using the log transform of Equation (6), the CUSUM test statistic is indeed simply given by:

$$W_n = \max(0, W_{n-1}) + \log(L_n), \quad n \geq 1, W_0 = 0 \quad (7)$$

By just computing a sum at each time sample and comparing W_n to a threshold, one can have a robust online detection technique. The question of estimating the change point ν can be solved easily by expanding Equation (5) as:

$$\begin{aligned}
\hat{\nu} &= \operatorname{argmax}_{0 \leq \nu < n} \prod_{k=\nu+1}^n L_k \\
&= \operatorname{argmax}_{0 \leq \nu < n} \sum_{k=\nu+1}^n \log(L_k), \\
&= \operatorname{argmax}_{0 \leq \nu < n} \sum_{k=1}^n \log(L_k) - \sum_{k'=1}^{\nu} \log(L_{k'}) \\
&= \operatorname{argmin}_{0 \leq k < n} (S_k).
\end{aligned} \tag{8}$$

with:

$$S_k = \sum_{i=1}^k \log \frac{f_1(x_i|\theta_1)}{f_0(x_i|\theta_0)}, \text{ and } S_0 = 0. \tag{9}$$

2.1.2. Example

To illustrate the CUSUM, let us consider a change of mean in a single data stream composed of independent Gaussian random variables:

$$\begin{aligned}
f(x_k|\theta_0) &= \mathcal{N}(x_k; \mu_0, \sigma^2) \\
f(x_k|\theta_1) &= \mathcal{N}(x_k; \mu_1, \sigma^2)
\end{aligned} \tag{10}$$

In this case the CUSUM test statistic can be computed as:

$$W_n = \max \left(0, W_{n-1} + \frac{\delta_\mu}{\sigma^2} \left(x_n - \mu_0 - \frac{\delta_\mu}{2} \right) \right) \tag{11}$$

where,

$$\delta_\mu = \mu_1 - \mu_0 \tag{12}$$




Figure 1: Example of detection of a change of mean in a single data stream from $\mu_0 = 0$ to $\mu_1 = 0.1$ with a Gaussian noise distribution of standard deviation $\sigma = 1$. The data stream is plotted at the top with the change-point time indicated by a vertical line. The CUSUM variable W_n is plotted on the middle figure with the value of the chosen threshold indicated by the horizontal red line and the time of detection indicated by the green vertical line. The variable from Equation (9) is plotted at the bottom with the estimated change point indicated by the vertical red line.

Figure 1 empirically shows that the CUSUM technique is able to detect a change in the mean of a signal which is not obvious by looking only at the time series x_n . The method has a delay to detect the change-point (51 time samples in this case) but it is able to detect it nonetheless. At the bottom of the figure we can see that $\arg \min(S_n)$, defined by Equation (9), gives us an estimate of the change point. This estimate can also be found at the last time $W_n = 0$. The CUSUM can be used to detect any changing parameter [17] even if it is most commonly used to detect a change of mean or variance. As in any detection technique there is a balance to make between detection rate and false alarm, here between the average detection delay and the ARL2FA.

A way to deal with the setting of the method is to determine what ARL2FA is tolerable, set the detection threshold (to which W is compared) to get the wanted ARL2FA and then check what average detection delay is obtained. When comparing several methods, one can set the thresholds so that the ARL2FA is the same among all of them and compare the average detection delays to determine which gives the quickest detection.

2.2. Finite moving average (FMA)

Concerned by the cases in which the change is temporary, [18] proposes a method to detect this change by computing a likelihood ratio test on a moving time window of the signal defined as:

$$Z_n^w = \sum_{k=n-w+1}^n S_k \quad (13)$$

which depends on the window length w . This test statistic, Z_n^w is then compared to a threshold to trigger or not a detection. In [18], the authors compared this approach to the CUSUM when the amplitude of a change in mean is lower than expected and when the change duration is finite or even if the change is intermittent. These characteristics are of the utmost interest for our purpose.

This technique is similar to the MOSUM (moving sum) [19] and requires to either memorise w values of the likelihood ratio S_k or to compute w times more operations every time sample than the CUSUM technique. Also, this method seems to be sensitive to the difference between the length of the window and the duration of the signal to detect.

2.3. Temporary-Event-CUSUM

In this paper we introduce a new technique called Temporary-Event-CUSUM (TE-CUSUM). Because the change is transitory, the model we consider is defined through the two hypotheses:

$$\begin{aligned} \mathcal{H}_0 : p(x_1, \dots, x_n) &= \prod_{k=1}^n f_0(x_k | \theta_0) \\ \mathcal{H}_1 : p(x_1, \dots, x_n; \nu, N) &= \prod_{k=1}^{\nu} f_0(x_k | \theta_0) \prod_{k=\nu+1}^N f_1(x_k | \theta_1) \prod_{k=N+1}^n f_0(x_k | \theta_0) \end{aligned} \quad (14)$$

Proposition 1. *The test statistic obtained by solving the generalised likelihood ratio test of Equation (14) can be recursively obtained as follows:*

$$G_n = \max(G_{n-1}, W_n), G_0 = 0, \quad (15)$$

with W_n being the CUSUM test statistic defined in Equation (7).

Proof. The GLRT of Equation (14) can be written as:

$$\begin{aligned} G_n &= \max_{[\nu, N]} \sum_{k=\nu}^{N \leq n} \log \frac{f_1(X_k | \theta_1)}{f_0(X_0 | \theta_0)} \\ &= \max_{0 < k \leq n} \left(\max_{\nu} \sum_{i=\nu}^k \log \frac{f_1(X_i | \theta_1)}{f_0(X_i | \theta_0)} \right) \\ &= \max_{0 \leq k < n} (W_k) \end{aligned} \quad (16)$$

which leads straightforwardly to the recursive form introduced in Prop. 1. \square

In Equation (16) it is implicit that ν is the last change-point before k . Moreover, causality forces $\hat{\nu} < n$ when W_n is computed.

On univariate cases, TE-CUSUM is strictly equivalent to the standard CUSUM because a test on G_n is equivalent to a test on W_n . From Equation (15), we have $G_n = \max_{0 \leq k \leq n} (W_k)$. Therefore if $W_k > G_{k-1}$ then $G_k = W_k$ if $W_k > h$ then $G_k > h$, h being the detection threshold. If $G_{k-1} < h$ and $W_k < G_{k-1}$ then $G_k = G_{k-1} < h$ and $W_k < G_k < h$. There is no way $W_k > h$ without $G_k > h$ and neither $G_k > h | G_{k-1} < h$ without $W_k > h$.

3. Multivariate detection

In this section, we consider a multi-sensor network which consists of a collection of indexed sensors $\mathcal{K} = \{1, \dots, L\}$ where each of them observes a realization from the previously discussed model. More specifically, under normal conditions, the distribution which governs the behaviour of each of the sensor is given by $f_0(\cdot | \theta_0)$. At random time and during some random duration, a change could occur which affects a subset of sensors $\mathcal{K}_c \in \{\emptyset \cup \mathcal{K}\}$. The detection problem of a change can be thus formulated using the following

binary hypothesis test model:

$$\begin{aligned}
\mathcal{H}_0 : p(\mathbf{x}_1, \dots, \mathbf{x}_n) &= \prod_{l \in \mathcal{K}} \prod_{k=1}^n f_0(x_{k,l} | \theta_0) \\
\mathcal{H}_1 : p(\mathbf{x}_1, \dots, \mathbf{x}_n; \boldsymbol{\nu}, \mathbf{N}) &= \prod_{l \in \mathcal{K}_c} \prod_{\substack{1 \leq k \leq \nu_l \\ N_l < k \leq n}} f_0(x_{k,l} | \theta_0) \prod_{k=\nu_l+1}^{N_l} f_1(x_{k,l} | \theta_1) \\
&\quad \times \prod_{j \in \mathcal{K} \setminus \mathcal{K}_c} \prod_{k=1}^n f_0(x_{k,j} | \theta_0)
\end{aligned} \tag{17}$$

3.1. A brief review of existing procedures

Solving the problem set by Equation (17) would necessitate to test all combination of change-points for each possible subset of sensors. Some optimisation approaches have been used by [10], or [12] and [20]. The major drawback of these approaches is that it becomes rapidly too computationally expensive and it loses the possibility of a recursive computation. [8] sees the multivariate CUSUM variable as the norm of the sum of the local test statistics and [14] takes into account the case where the size of the subset of sensors is roughly known. In some cases, when the number of sensors is very large [13] proposed to only merge binary units to the decision center so that detection is triggered by their number and not by an aggregation of local values.

Except for [12] and [20], all the developed methods consider that the change is permanent ($N_l = \infty$) for all the affected sensors. The Multi-dimension exponentially weighted moving average (MD-EWMA) algorithm presented by [21] takes into account temporary changes but requires the change to be synchronous between the sensors.

In this work, we intend to remove the limitation $N_l = \infty$ (thus allowing some sensors to stop monitoring the change at some point) without increasing the computational power required for the detection.

Two basic ways to adapt the CUSUM to multivariate cases is by computing the sum of local variables or by extracting the maximum value among the local CUSUM variables.

The SumCUSUM [9] associates local CUSUM variables as follows:

$$T_{SC}(n) = \frac{1}{L} \sum_{l=1}^L W_{l,n}. \tag{18}$$

Where $T_{SC}(n)$ being the global SumCUSUM variable; *i.e.* the sum of the local CUSUM variables of the L data stream. $W_{l,n}$ being the CUSUM variable of the l^{th} sensor at time n . This variable $T_{SC}(n)$ will be compared to an adapted threshold h to make a decision of a detection when $T_{SC}(n) > h$.

The MaxCUSUM extracts the highest value among the local CUSUM values as shown by Equation (19):

$$T_{MC}(n) = \max_{l \in [1, \dots, L]} (W_{l,n}). \quad (19)$$

It appears that SumCUSUM will be relevant to be used when all or almost all of the data streams are affected by the signal while MaxCUSUM will be relevant when one or only a few of the data streams are affected.

[9] proposed to select ("censor") sensors and compute a partial and optimized SumCUSUM with a low computational cost. It seems to be a very effective way to merge the data for an online use of the method. The SumCUSUM variable is thus transformed as:

$$T_{cSC}(n) = \frac{1}{\sum \mathbb{1}_{W_{l,n} \geq c}} \sum_l^L W_{l,n} \mathbb{1}_{W_{l,n} \geq c}. \quad (20)$$

with c a threshold based on the prior rough knowledge of the value $W_{l,n}$ would take if it were affected by the signal. Figure 2 shows the results of the three methods depending on the proportion of sensors affected by the signal. The average run-length to false alarm (ARL2FA) of all three methods have been set to 30 000. We can see that when 1 or 2 out of 10 sensors are affected, the MaxCUSUM shows lower detection delays. When 3 or more sensors are affected, the SumCUSUM gives a quicker detection than the MaxCUSUM. We can also infer from Figure 2 that the censored SumCUSUM is a good compromise between SumCUSUM and MaxCUSUM. However, [9] shows that the best choice for c depends on the number of sensors affected. While in some cases the proportion of sensors affected can be roughly predicted, in most cases it is completely unknown. In the case of the example of Figure 2 c has been set as 60% of the global threshold h .




Figure 2: Average detection delay of SumCUSUM, MaxCUSUM and censored SumCUSUM techniques on a change in the mean of a Gaussian distribution with signal to noise ratio of -6dB. Change-point occurs at time sample 1000.

3.2. A novel adaptive censoring technique

To overcome the limitation of requiring some prior knowledge on the expected values of $W_{l,n}$ in order to carefully choose the absolute threshold c from [9], we propose a relative threshold c_n , computed for every time sample by:

$$c_n = \alpha \times \max_{l \in [1, \dots, L]} (W_{l,n}) \quad (21)$$

with α being a factor so that $0 \leq \alpha \leq 1$.

In both cases the censoring technique is a compromise between the SumCUSUM and the MaxCUSUM. The results given by the two can be retrieved using particular values for c (0 and h) or α (0 and 1).

In order to assess the difference of behaviour of the two threshold types regarding the number of affected sensors when it is unknown, we conducted an experiment which results are shown in Figure 3. In this experiment, h is set such that an average run-length to false alarm (ARL2FA) of 10,000

is obtained and c and α values are set to be those which give the overall quickest detection for an unknown number of affected sensors between 1 and 20. The results empirically show that the proposed adaptive censoring technique outperforms the classical one. The gap in performance increases with the number of affected sensors.




Figure 3: Average detection delay standard and adaptive Censored SumCUSUM at -12dB (with optimized values: $c = 5.8$ and $\alpha = 0.6$). A change of mean in the Gaussian distributions of a subset (in abscissa) of sensors appears at time 1000 while the other sensors keep the centered Gaussian distribution.

Because we cannot know in advance the number of sensors that will be affected in addition of the non-requirement of the knowledge of the expected values of $W_{l,n}$, the relative threshold is consequently more relevant.

Some clues can be pointed out to explain this difference of behaviour by examining the differences between the two methods in some particular cases.

Case 1: All the local $W_{l,n} > c$ have close values one to another and are relatively far from the $W_{l,n} < c$. In this case both methods will compute the same T_{cSC} .

Case 2: The local $W_{l,n} > c$ have very different values. The standard method computes T_{cSC} adding more low values of $W_{l,n}$ and has a lower value of T_{cSC} which slows the detection.

Case 3: All the local $W_{l,n} > c$ have close values one to another and the $W_{l,n} < c$ are also close to the others. In this case, it is the standard method that computes a highest value for T_{cSC} , but the case implies that all the values are close to c so detection does not happen in both cases unless the value chosen by c is close to h and in this case we have a behaviour close to the MaxCUSUM.

In all that follows we apply this optimised relative censoring technique to all local statistics (CUSUM, TC-CUSUM, FMA) and keep the SumCUSUM and the MaxCUSUM as benchmarks.

3.3. Asynchronous monitoring and Temporary-Event-CUSUM on multivariate cases

In the previous section we have considered that the signal appears simultaneously on all the affected sensors. Indeed, all the local test variables are computed simultaneously and it is from these that we can compute the global variable at time n and make a decision regarding the detection.

In many practical cases the signal can be monitored by the sensors with a delay. Even more, some sensors can cease to be affected by the signal before some others begin to be. Thus, the sensors are not affected at the same time. One could say that we should try to find the best synchronicity of the data streams, meaning the synchronicity which maximises the associated CUSUM variable but this is a combinatorial problem.

By using locally the novel TE-CUSUM test statistic defined in Proposition 1, the Sum-TE-CUSUM allows us to get the best synchronicity without requiring to test all the combinations, and thus saving a lot of computational resources. This time the global test variable becomes:

$$T_{STEC}(n) = \frac{1}{L} \sum_{l=1}^L \max_{[\nu_l, N_l]} \sum_{k=\nu_l}^{N_l \leq n} \log \frac{f_{1,l}(X_{k,l} | \theta_{1,l})}{f_{0,l}(X_{k,l} | \theta_{0,l})} \quad (22)$$

ν_l is the change-point for the l^{th} data stream and N_l is the end of the signal presence in the l^{th} data stream.

As a reminder of Equation (15), the local variable is:

$$G_n = \max(G_{n-1}, W_n) = \max_{0 < k \leq n} (W_k) \quad (23)$$

From Equation (16) and with (ν, N) of the l^{th} data stream being rewritten (ν_l, N_l) :

$$T_{STEC}(n) = \frac{1}{L} \sum_{l=1}^L G_{n,l} \quad (24)$$

The censoring technique can be applied to Equation (24) simply by adding a threshold like in Equation (20):

$$T_{cSTEC}(n) = \frac{1}{\sum \mathbb{1}_{G_{n,l} > c_n}} \sum_{l=1}^L G_{n,l} \times \mathbb{1}_{G_{n,l} > c_n} \quad (25)$$

Here is an example to illustrate the Sum-TE-CUSUM (T_{STEC}). An event is monitored in three data stream but with such a delay that there is no overlap. In Figure 4, we can see the three data streams with and without noise.




Figure 4: Three data streams monitoring a short event without overlapping (noiseless in red, noise+signal in blue)

Figure 5 shows the evolution of the test variable of the standard Sum-CUSUM technique and the TE-CUSUM.




Figure 5: Evolution of test variables on the data streams of Figure 4

We can see in Figure 5 that the SumCUSUM decreases between each appearance of the signal while the TE-CUSUM stands by and increases again as soon as the signal appears on another data stream. With the TE-CUSUM, we can detect the presence of the event with a higher threshold. If we set both thresholds in order to have a probability of false alarm of 1% on this interval and if we make 10,000 runs we obtain a detection rate of 41% with the SumCUSUM and 83% with TE-CUSUM. To achieve 1% of false alarm, detection thresholds are set to 17.5 for the TE-CUSUM and to 13 for the SumCUSUM.

A pseudo-code, explaining the TE-CUSUM function is available in Appendix A.

Remark: In order to compare it to the other methods, we can also use the censoring technique to extend the FMA technique to multivariate cases:

$$T_{cFMA}(n) = \frac{1}{\sum_l \mathbb{1}_{Z_{n,l}^w > c_n}} \sum_{l=1}^L Z_{n,l}^w \times \mathbb{1}_{Z_{n,l}^w > c_n} \quad (26)$$

4. Validation

In this section we compare the different detection methods presented previously. The studied methods are used to detect a change in the mean amplitude A affecting only a subset of sensors. The measurement noise at each sensor is assumed to be normally distributed with zero mean and standard deviation σ . This experiment is conducted on several cases with two different signal to noise ratios (SNR) defined as $SNR = 10 \log_{10}(\frac{A^2}{\sigma^2})$. We ensured to be out of transitional mode in the signal by exposing the first affected sensor after 1500 time samples. In the following plots for simplicity, zero corresponds to the beginning of exposure and therefore not the beginning of the experiment. Ten sensors are considered among which 3 or 7 monitor the event. The case where 5 sensors monitor the signal is displayed in Appendix B. The censoring technique will be applied to TE-CUSUM and FMA. For each method the global threshold h is set to have an ARL2FA of 30,000 time samples, and the α parameter from the censoring technique is set (except for SumCUSUM and MaxCUSUM) for the quickest detection for a random number of signal but when all of them are synchronized. In order to set the ARL2FA to 30,000 time samples, the detection thresholds for each method have been set so we have the cumulative probabilities of false alarm presented Figure 6

Two sizes of window have been chosen for the FMA technique. One window with 50 time samples and another with 200 time samples. Because the duration of the event is supposed to be unknown, this will show a case where the time window is longer than the exposure and a case where the time window shorter than the exposure. In the scenario where the SNR is $-8dB$. In the second scenario, where the SNR is $-14dB$, the duration of the exposure will, fortunately for the FMA200, be of 200 time samples.

The numerical experiments are divided into four cases: the first case is when all the signals are monitored simultaneously, in the second case there is a drift of half the signal length between each sensor which monitors the signal, in the third case the drift is of a full signal length (This case is displayed in appendix Appendix B) and in the fourth the drift is of one and a half signal length. This last simulation tends to represent a diffuse event passing through a field of sensors so that these ones do not monitor the event at the same time and only a little part of the sensors can monitor the signal anyways.





Figure 6: Empirical cumulative false alarm probability for the experiment presented in the validation section. All the methods have been set to have close values of false alarm probability. From this figure we can deduce that the probability of a false alarm to happen in a 1000 time samples window is approximately of 5%. This means that there is a 5% chance of false alarm before the beginning of the signal, 5% chance that a false alarm leads to detection in the first scenario and 10% chance that a false alarm leads to detection in the second scenario. (SC=SumCUSUM, MC=MaxCUSUM, cSC = Censored-SumCUSUM, cSTEC = Censored-Sum-Temporary-Event-CUSUM, cFMA50 = Censored-Finite-Moving-Average with a window of 50 time samples and cFMA200 = Censored-Finite-Moving-Average with a window of 200 time samples).

4.1. First scenario

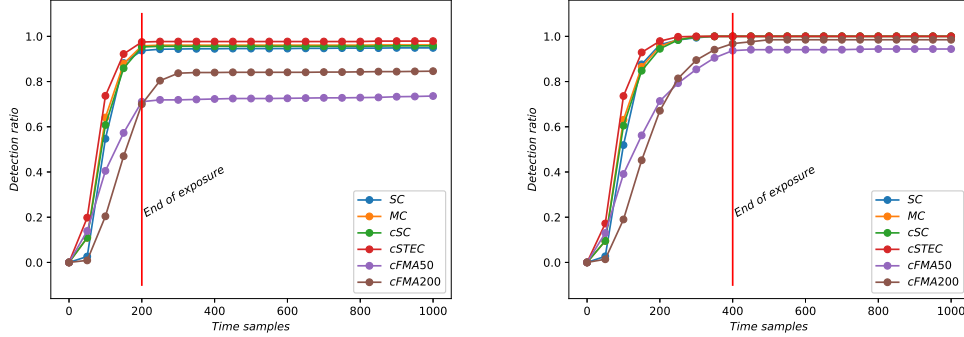
In this first scenario, all the studied techniques are set to expect an offset of 0.4 in amplitude and to have an ARL2FA of 30,000 time samples. When a sensor is affected by the signal, its mean value is affected by an offset of 0.4 for a duration of 100 time samples. We expect the detection ratio of the methods to grow rapidly and then become steady shortly after the end of exposure as no signal is added to the system anymore. The growth of each slope will indicate how fast the detection is, and the value reached after the end of exposure indicates the chances of detection given the current experiment.



(a) Results when 3 sensors out of 10 are affected (b) Results when 7 sensors out of 10 are affected

Figure 7: Cumulative detection rate as a function of time of each of the presented method when streams are synchronized with a signal to noise ratio of -8dB. The end of exposure marks the time until at least one sensor is monitoring the change in amplitude. The exposure time is of 100 time samples.(Legend is detailed in Figure 6).


From Figures 7a and 7b, we can see that when all the exposures are synchronized, the SumCUSUM technique is the one that gives the best results. However, the censored-SUM-TE-CUSUM and the Max-CUSUM give rather good results in this case. When data streams stop to monitor the event, at the time marked "end of exposure", the detection ratio is over 80% for the four best methods in the case where 3 data sensors are exposed and over 90% when 7 sensors are exposed. In the case 7 sensors are exposed the censored TE-CUSUM as well as the SumCUSUM give a result of almost 100% of detection at the end of exposure. The FMA technique seems to give slower detection and does not manage to reach the detection rate of the other techniques.



(a) Results when 3 sensors out of 10 are affected (b) Results when 7 sensors out of 10 are affected

Figure 8: Cumulative detection rate as a function of time of each of the presented method when streams are slightly out-of-sync with a signal to noise ratio of -8dB. The end of exposure marks the time until at least one sensor is monitoring the change in amplitude. Here the delay between the beginning of the exposure of a sensor and the next is of 50 time samples while the exposure time by sensor is of 100 time samples.(Legend is detailed in Figure 6).

The results presented on Figures 8a and 8b are obtained in the same conditions except this time, the start of monitoring of every data stream is delayed by 50 time samples from the previous one. Here, the censored TE-CUSUM shows how it manages to give similar results in more difficult conditions. Indeed, at the end of the exposure in figure 8a the total energy transmitted by the event to the system is the same than at the end of exposure in figure 7a. The censored TE-CUSUM gives in both cases a 90% detection rate at the "end of exposure" where the SumCUSUM technique falls from about 95% when signals are synchronized to a little less than 90% when they are slightly out-of-sync. On these figures we can also see that the longer the overall exposure is the better the FMA results seem to be.



(a) Results when 3 sensors out of 10 are affected (b) Results when 7 sensors out of 10 are affected

Figure 9: Cumulative detection rate as a function of time of each of the presented method when streams are totally out-of-sync with a signal to noise ratio of -8dB. The end of exposure marks the time until at least one sensor is monitoring the change in amplitude. In (b) the end of exposure happen after the end of monitoring. Here the delay between the beginning of the exposure of a sensor and the next is of 150 time samples while the exposure time by sensor is of 100 time samples. (Legend is detailed in Figure 6).


In the case of Figures 9a and 9b there is a gap of 50 time samples where no data stream monitors the signal between the exposure of each data stream. One can see that the detection rate tends to stay still between the exposure of each data stream. In figure 9b the end of exposure happens at time sample 1100 and is not marked on the figure. In these two last cases we can see that the censored FMA technique, can give good results when the exposure lasts. It is notable that in all the non synchronized cases, censored TE-CUSUM gives the best early detection and the best detection rate. This time, the SumCUSUM, because it only considers instantaneous CUSUM variables, has a very low detection rate compared to the other techniques. The difference of performances between the SumCUSUM and its censored version clearly shows the benefit of using our proposed censoring technique when signals are not synchronised. A similar behaviour has been observed with the censoring technique applied to our TE-CUSUM.

Results with more cases are presented in Appendix B.

4.2. Second scenario

In this second scenario, we run the same tests, but with a lower SNR. The amplitude of the signal expected by the methods is still of 0.4, but the real signal has an amplitude of only 0.2. Again, detection thresholds of all


methods are set to have an ARL2FA of 30,000 time samples. When a sensor is affected by the signal, its mean value is affected by an offset of 0.2 for a duration of 200 time samples.



(a) Results when 3 sensors out of 10 are affected (b) Results when 7 sensors out of 10 are affected

Figure 10: Cumulative detection rate as a function of time of each of the presented method when streams are synchronized with a signal to noise ratio of -14dB. The end of exposure marks the time until at least one sensor is monitoring the change in amplitude. The exposure time is of 100 time samples for the affected sensors. (Legend is detailed in Figure 6).

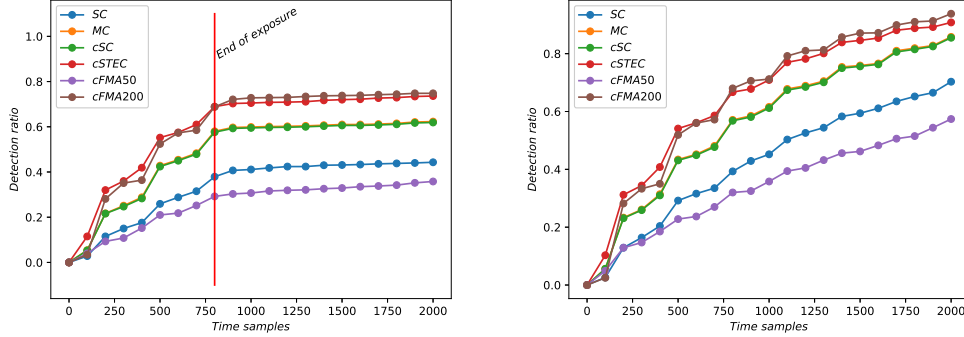
Figures 10a and 10b shows that, this time, nearly all the methods have the utmost difficulties to give high detection rates. The exposure of more data streams in the case of Figure 10b improves the detection rates of the methods, but SumCUSUM is the only one which is fully satisfying.



(a) Results when 3 sensors out of 10 are affected (b) Results when 7 sensors out of 10 are affected

Figure 11: Cumulative detection rate as a function of time of each of the presented method when streams are slightly out-of-sync with a signal to noise ratio of -14dB . The end of exposure marks the time until at least one sensor is monitoring the change in amplitude. Here the delay between the beginning of the exposure of a sensor and the next is of 50 time samples while the exposure time by sensor is of 100 time samples. (Legend is detailed in Figure 6).

Figures 11a and 11b show that when signals are slightly out-of-sync, the censored FMA (with a time window of 200) and censored TE-CUSUM give similar results. If these results are not fully satisfying, it is important to remember that the SNR is very low. In terms of power, the signal to noise ratio is of -14dB on a data stream when the signal is present.



(a) Results when 3 sensors out of 10 are affected (b) Results when 7 sensors out of 10 are affected

Figure 12: Cumulative detection rate as a function of time of each of the presented method when streams are totally out-of-sync with a signal to noise ratio of -14dB. The end of exposure marks the time until at least one sensor is monitoring the change in amplitude. In (b) the end of exposure happen after the end of monitoring. Here the delay between the beginning of the exposure of a sensor and the next is of 150 time samples while the exposure time by sensor is of 100 time samples. (Legend is detailed in Figure 6).

In this last example, Figures 12a and 12b show that the censored TE-CUSUM and the censored FMA can reach the same detection rate as when signals are synchronized or slightly out-of-sync, even if these detection rate are reached later than when signals are synchronised. The SumCUSUM gives very good results when most sensors monitor the event and when these are synchronized. However, when one of these condition is not met, the SumCUSUM performance decreases very fast.

Results with subsets of affected sensors and one additional asynchronous case are presented in Appendix B.

5. Conclusion

In this paper, we have addressed the detection problem of an event which only appears on a subset of sensors, and such that these appearances can be delayed one to another so that they can be perceived by the system as if the data streams monitoring the same event were out-of-sync.

Existing methods have already explored the fact that an event can be monitored by only a portion of the sensors, as well as they can deal with the fact that the change point does not occur at the same time on every data stream. But what standard CUSUM methods lack to consider is the fact

that a sensor can in some cases cease to monitor the event while another one does.

The method we propose takes in consideration all the cases of delays between the change point on the different data streams as well as the fact that some of them can cease to monitor the event before the end of the system exposure. We have shown that if the system is composed of only one data stream, the TE-CUSUM is equivalent to a standard CUSUM procedure.

We have also shown that the censored TE-CUSUM, beside the cases where all signals are synchronized, gives the best results, even if those can be similar to the FMAs at very low SNR when the FMA window is adapted to the signal length. It is important however to note that the FMA was not originally designed to be used on multivariate out-of-sync cases. Also, a big advantage of the censored TE-CUSUM, is that it keeps the recursive computation of the CUSUM. Indeed, it only adds a comparison of the SumCUSUM variable to the last maximum to the SumCUSUM technique where the FMA computes the likelihood ratio on a signal portion which can be rather long, which means it needs to store many values (Appendix C) to compute the test variable. The FMA also requires to use the window length which can be sensitive to the length of the signal that is expected. This sensitive parameter of the window length for the FMA gives the TE-CUSUM the advantage of being easier to tune when the length of exposure is not known.

In this paper we have introduced the new TE-CUSUM method which provides a light and simple detection technique which covers a greater range of cases by adding the possibility for the event to detect to be temporarily and not simultaneously monitored by the different data streams while still giving rather good results in standard cases when data streams are synchronised. The proposed procedure can have many practical applications, e.g. when a network of sensors is monitoring a localised event passing through. It can be, for instance a plume travelling into the air containing a chemical compound one wishes to detect or a furtive object passing through several radar monitored areas one after the other.

Appendix A. TE-CUSUM pseudo-code

```

Detection  $\leftarrow$  False
t  $\leftarrow$  0
for  $l \in L$  do
     $G_{l,0} \leftarrow 0$ 
     $W_{l,0} \leftarrow 0$ 
end for
while Detection = False do
     $t \leftarrow t + 1$ 
    for  $l \in L$  do
         $X_{l,t} \leftarrow$  value of sensor l at time t
         $L_{l,t} \leftarrow$  log - likelihood for sensor l
        if  $W_{l,t-1} \geq 0$  then
             $W_{l,t} \leftarrow W_{l,t-1} + L_{l,t}$ 
        else
             $W_{l,t} \leftarrow L_{l,t}$ 
        end if
        if  $W_{l,t} > G_{l,t-1}$  then
             $G_{l,t} \leftarrow W_{l,t}$ 
        else
             $G_{l,t} \leftarrow G_{l,t-1}$ 
        end if
    end for
     $TECUSUM \leftarrow 0$ 
     $s \leftarrow 0$ 
    for  $l \in L$  do
        if  $G_{l,t} > \alpha \times \max_l G_{l,t}$  then
             $TECUSUM \leftarrow TECUSUM + G_{l,t}$ 
             $s \leftarrow s + 1$ 
        end if
    end for
    if  $\frac{TECUSUM}{L_t} \geq$  threshold then
        Detection  $\leftarrow$  True
    end if
end while

```

Appendix B. Experimental results

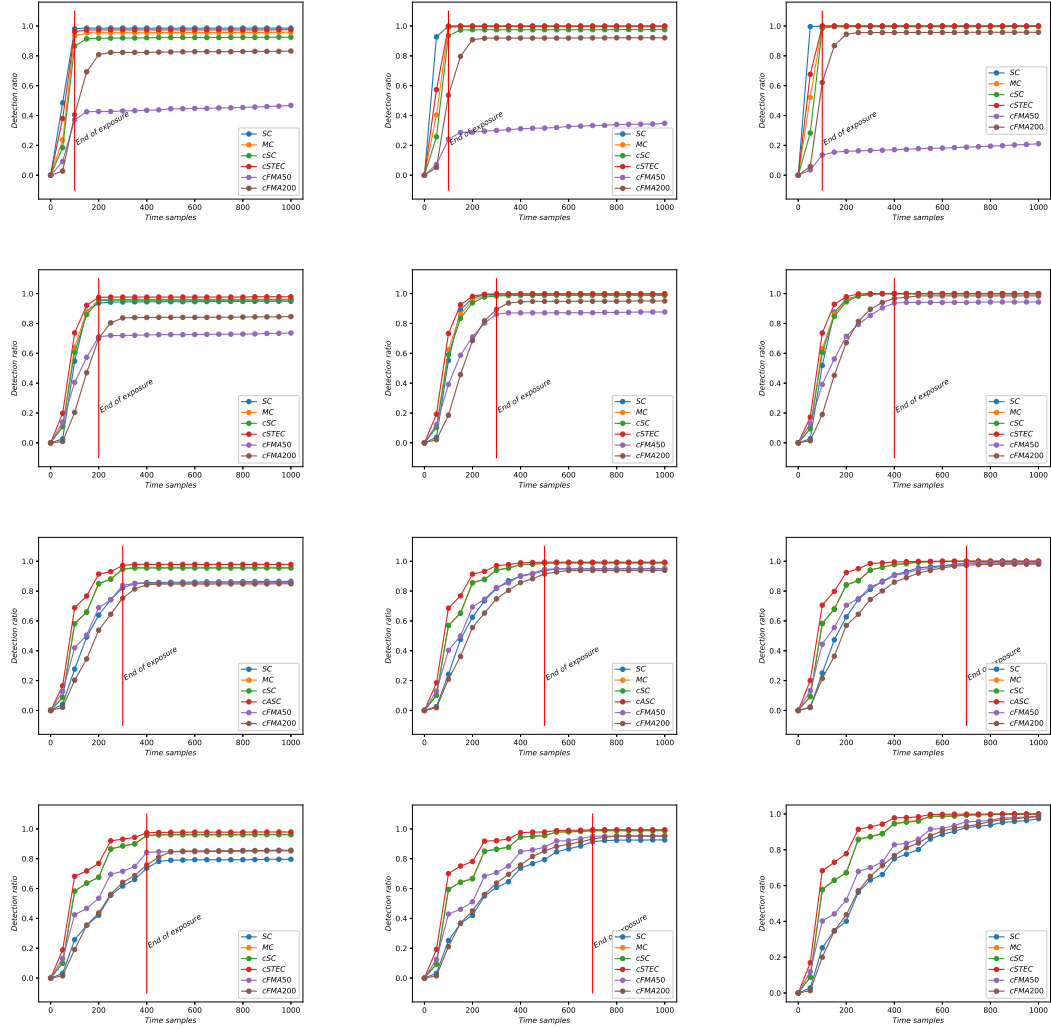


Table B.1: Results with 3/10 affected data streams on the first column, 5/10 on the second column and 7/10 on the third. With synchronised signals on the first row, a 50 time samples delay on the second, a full signal length delay on the third, and a 50 sample gap between two exposure on the fourth row. Signal amplitude is 0.4 and there is 100 time samples by exposure.




Table B.2: Results with 3/10 affected data streams on the first column, 5/10 on the second column and 7/10 on the third. With synchronised signals on the first row, a 100 time samples delay on the second, a full signal length delay on the third, and a 100 sample gap between two exposure on the fourth row. Signal amplitude is 0.2 and there is 200 time samples by exposure.

Appendix C. Computational resources

method	description	number of computations	number of stored variables
MaxCUSUM	L access g + L tests	2L	L
SumCUSUM	L access g + L sums	2L	L
CensoredSC	L access g + L test + < L sums	< 3L	L
TE-CUSUM	same as cSC + L tests on G	< 5L	2L
FMA	$L \times 2$ (access + sums) + L tests + < L sums	< 6L	$w \times L$

Table C.3: Estimation of the computational cost of each method (L is the number of sensors and w is the window length for the FMA). In our cases $L \ll w$. Here the computation of the likelihood ratio is not displayed as it is common to all the techniques. The global decision test is not displayed for the same reason. Computer operations have been categorised in 3 type of operation; memory access to local values of g, tests between two variables and sum. To simplify, every operation listed here are considered to have the same computational cost.

As an example, we can describe the functioning of the TE-CUSUM based on the algorithm pseudo-code of Appendix A:

Dots 1, 2 and 7 are common to all the techniques so these are not displayed in the comparison table.

Dots 1, 2, 3, 4 and 7th are common to all the CUSUM based techniques.

Dot 6 is common to Censored Sum-CUSUM and Censored TE-CUSUM.

At every measure time:

- L values are recorded by the sensors.
- for each of these L values, the log-likelihood is computed
- L values of W are loaded from the memory
- L tests between W and $W + L$ are performed to update W
- L tests between G and W are performed to update G
- $l < L$ Sum are performed to get the global test statistic
- the global test statistic is compared to a threshold to trigger or not detection

References

- [1] H. Rajaona, F. Septier, P. Armand, Y. Delignon, C. Olry, A. Albergel, J. Moussafir, An adaptive bayesian inference algorithm to estimate the parameters of a hazardous atmospheric release, *Atmospheric Environment* 122 (2015) 748–762. doi:10.1016/j.atmosenv.2015.10.026.
- [2] F. Septier, P. Armand, C. Duchenne, A bayesian inference procedure based on inverse dispersion modelling for source term estimation in built-up environments, *Atmospheric Environment* 242 (2020) 117733. doi:https://doi.org/10.1016/j.atmosenv.2020.117733.
- [3] T. D. Popescu, Signal segmentation using changing regression models with application in seismic engineering, *Digital Signal Processing* 24 (2014) 14–26. doi:https://doi.org/10.1016/j.dsp.2013.09.003.
- [4] P. Braca, D. Gaglione, S. Marano, L. M. Millefiori, P. Willett, K. R. Pattipati, Quickest detection of covid-19 pandemic onset, *IEEE Signal Processing Letters* 28 (2021) 683–687. doi:10.1109/LSP.2021.3068072.
- [5] A. Aue, L. Horváth, Structural breaks in time series, *Journal of Time Series Analysis* 34 (1) (2013) 1–16. arXiv:https://onlinelibrary.wiley.com/doi/pdf/10.1111/j.1467-9892.2012.00819.x, doi:https://doi.org/10.1111/j.1467-9892.2012.00819.x.
- [6] M. Shbat, V. Tuzlukov, Primary signal detection algorithms for spectrum sensing at low snr over fading channels in cognitive radio, *Digital Signal Processing* 93 (2019) 187–207. doi:https://doi.org/10.1016/j.dsp.2019.07.016.
- [7] E. S. Page, Continuous inspection schemes, *Biometrika* 41 (1-2) (1954) 100–115. arXiv:https://academic.oup.com/biomet/article-pdf/41/1-2/100/1243987/41-1-2-100.pdf, doi:10.1093/biomet/41.1-2.100.
- [8] V. Golosnoy, S. Ragulin, W. Schmid, Multivariate cusum chart: properties and enhancements, *AStA Advances in Statistical Analysis* 93 (3) (2009) 263–279.
- [9] Y. Mei, Efficient scalable schemes for monitoring a large number of data streams, *Biometrika* 97 (2) (2010)

419–433. arXiv:<https://academic.oup.com/biomet/article-pdf/97/2/419/583997/asq010.pdf>, doi:10.1093/biomet/asq010.

- [10] M. N. Kurt, X. Wang, Multisensor sequential change detection with unknown change propagation pattern, *IEEE Transactions on Aerospace and Electronic Systems* 55 (3) (2019) 1498–1518. doi:10.1109/TAES.2018.2873067.
- [11] A. Tartakovsky, I. Nikiforov, M. Basseville, *Sequential Analysis: Hypothesis Testing and Change-point Detection* p. 111, 2014.
- [12] G. Rovatsos, V. V. Veeravalli, D. Towsley, A. Swami, Quickest detection of growing dynamic anomalies in networks (2020). arXiv:1910.09151.
- [13] T. Banerjee, V. V. Veeravalli, Data-efficient quickest change detection in sensor networks, *IEEE Transactions on Signal Processing* 63 (14) (2015) 3727–3735. doi:10.1109/TSP.2015.2432737.
- [14] B. Y. Xie, D. Siegmund, Sequential multi-sensor change-point detection, in: *2013 Information Theory and Applications Workshop (ITA)*, 2013, pp. 1–20. doi:10.1109/ITA.2013.6502987.
- [15] H. L. V. Trees, *Detection, Estimation, and Modulation Theory: Radar-Sonar Signal Processing and Gaussian Signals in Noise*, Krieger Publishing Co., Inc., USA, 1992.
- [16] S. Knoth, Steady-state average run length(s): Methodology, formulas, and numerics, *Sequential Analysis* 40 (3) (2021) 405–426. arXiv:<https://doi.org/10.1080/07474946.2021.1940501>, doi:10.1080/07474946.2021.1940501. URL <https://doi.org/10.1080/07474946.2021.1940501>
- [17] S. Lee, J. Ha, O. Na, S. Na, The cusum test for parameter change in time series models, *Scandinavian Journal of Statistics* 30 (4) (2003) 781–796. arXiv:<https://onlinelibrary.wiley.com/doi/pdf/10.1111/1467-9469.00364>, doi:<https://doi.org/10.1111/1467-9469.00364>.
- [18] A. G. Tartakovsky, N. R. Berenkov, A. E. Kolessa, I. V. Nikiforov, Optimal sequential detection of signals with unknown appearance and disappearance points in time, *IEEE Transactions on Signal Processing* 69 (2021) 2653–2662. doi:10.1109/tsp.2021.3071016.

- [19] J. Noonan, A. Zhigljavsky, Power of the mosum test for online detection of a transient change in mean, *Sequential Analysis* 39 (2) (2020) 269–293. arXiv:<https://doi.org/10.1080/07474946.2020.1767406>, doi:10.1080/07474946.2020.1767406.
URL <https://doi.org/10.1080/07474946.2020.1767406>
- [20] G. Rovatsos, S. Zou, V. V. Veeravalli, Sequential algorithms for moving anomaly detection in networks, *Sequential Analysis* 39 (1) (2020) 6–31. arXiv:<https://doi.org/10.1080/07474946.2020.1726678>, doi:10.1080/07474946.2020.1726678.
- [21] Y. Wu, W. B. Wu, Sequential detection of common transient signals in high dimensional data stream, *Naval Research logistics* (2021). arXiv:<https://doi.org/10.1002/nav.22034>.
URL <https://doi.org/10.1002/nav.22034>