Restarted Bayesian Online Change-point Detector achieves Optimal Detection Delay

Réda Alami [1] **Odalric Maillard** [2] **Raphael Féraud** [3]

Abstract

In this paper, we consider the problem of sequential change-point detection where both the change-points and the distributions before and after the change are assumed to be unknown. For this problem of primary importance in statistical and sequential learning theory, we derive a variant of the Bayesian Online Change Point Detector proposed by (Fearnhead & Liu, 2007) which is easier to analyze than the original version while keeping its powerful message-passing algorithm. We provide a non-asymptotic analysis of the false-alarm rate and the detection delay that matches the existing lower-bound. We further provide the first explicit high-probability control of the detection delay for such approach. Experiments on synthetic and real-world data show that this proposal outperforms the state-of-art change-point detection strategy, namely the Improved Generalized Likelihood Ratio (Improved GLR) while compares favorably with the original Bayesian Online Change Point Detection strategy.

1. Introduction and related works

The problem of online detecting abrupt variations (change-points) in the generative parameters of a sequence of observations $x_1, \ldots x_n$, where observations are received one by one, is considered. Addressing this problem is useful in a number of real-world applications including finance (Schmitt et al., 2013), genetics (Grzegorczyk & Husmeier, 2009), cybersecurity (Polunchenko et al., 2012), robotics (Goldberg & Matarić, 2003; Biswas et al., 2002; Konidaris et al., 2010), speech recognition (Panda & Nayak, 2016), climate modeling (Nandhini & Devasena, 2019). The online change-point detection problem has received a lot of attention from various areas of mathematical statistics, information theory and computer science over the past century. We refer the interested reader to the recent survey (Aminikhanghahi & Cook, 2017) on the large amount of methods developed for time series change point detection, and to (Basseville et al., 1993; Brodsky & Darkhovsky, 1993; Jie & Gupta, 2000; Tartakovsky, 1991; Csörgö & Horváth, 1997; Wu, 2007) for classical textbooks on change-points. As noticed in (Aminikhanghahi & Cook, 2017), performance guarantees are still lacking for many such methods, especially in terms of finite time guarantee on the detection delay and estimation of the change-gap, both important features for the practitioner. Amongst the many methods, the celebrated CUSUM strategy from (Page, 1954) and its extension called Generalized Likelihood Ratio (GLR), that are following a frequentist approach based on likelihood ratio thresholding have been analyzed recently first in (Lai & Xing, 2010) and then in (Maillard, 2019), where a fully explicit parameter tuning is also provided, together with fully non-asymptotic guarantees.

In this paper, we turn to Bayesian approaches. In the seminal paper of (Fearnhead & Liu, 2007), the authors introduced the Bayesian Online Changepoint Detection (BOCPD) strategy to infer the most recent change-point, by computing the probability distribution of the elapsed time since the last change-point (runlength). Although the algorithm has been used extensively (including in non-stationary multi-armed bandits, (Mellor & Shapiro, 2013; Alami et al., 2016; Kerkouche et al., 2018; Dakdouk et al., 2018) and other change-point context (Adams & MacKay, 2007; Turner et al., 2009; Xuan & Murphy, 2007; Wilson et al., 2010; Saatçi et al., 2010; Caron et al., 2012; Niekum et al., 2014; Turner et al., 2013; Ruggieri & Antonellis, 2016; Knoblauch & Damoulas, 2018; Knoblauch et al., 2018)), up to our knowledge, no formal analysis of its performance in terms of false-alarm or detection delay has been performed except the work in (Knoblauch et al., 2018), where the authors has built a robust BOCPD version to reduce false discovery rates.

Note that although BOCPD stands for Bayesian Online Change Point Detection, the algorithm performs no detection at all; rather, it maintains weights to estimate the elapsed time since the last change-point. Following this work, we

[1]Total S.A. [2]INRIA-SCOOL Team [3]Orange Labs. Correspondence to: <reda.alami@total.com>, <alami1859@hotmail.fr>.

Proceedings of the 37^{th} International Conference on Machine Learning, Online, PMLR 119, 2020.

provide a modification of the BOCPD strategy that we analyze. In particular we provide a non-asymptotic guarantees related to the false-alarm (that is, detecting a change point while there was no change) in Theorem 2 and related to the detection delay (the number of steps after a change-point occurs before we declare detection) in Theorem 3.

In Section 2, we formally introduce the times-series model with abrupt changes, as well as notations. We provide in Section 3 a new formulation of the BOCPD strategy from (Adams & MacKay, 2007), that we reinterpret from the standpoint of aggregation of forecasters, leading to a compact formulation presented in Algorithm 1. We then present a simple way to make use of this strategy to effectively detect changes, instead of just estimating the time since the last change. We note that the analysis of BOCPD involves dealing with a combinatorial number of terms, and propose a simplification of this strategy in order to derive performance guarantees. We call the resulting strategy R-BOCPD for *Restarted Bayesian Online Change Point Detection*. Then, we provide in Section 5 the two theoretical guarantees of this strategy: namely the false alarm rate control and detection delay (Theorem 2, Theorem 3). Finally, we show numerically that this strategy outperforms its previous version BOCPD and its compares favorably with the Improved GLR strategy introduced by (Maillard, 2019). For the sake of clarity, the proofs of the analytical results are in the appendices.

2. Sequential change-point detection setting

Sequential change-point detection, which is rooted in classical statistical sequential analysis (Basseville et al., 1993), aims to detect the change in underlying distributions of a sequence of observations as quickly as possible.

In this paper, we study the online change point detection problem, where a sequence of independent univariate random variables with common fluctuation upper bound are collected, and the mean may change at one or multiple time points. Indeed, we consider an agent aiming at detecting changes in the generation of an online stream. At each time step t, the agent observes the datum $x_t \sim \mathcal{B}(\mu_t)$: a random variable following the Bernoulli distribution of mean μ_t and need to decide whether or not there is a change in the generation of the stream. Alternatively, the agent may compute at each time step t, an estimation $\widehat{\tau}_t$ of the last change-point.

Definition 1 (Piece-wise stationary Bernoulli process)**.** *Let T denote the time horizon of the game (stream length) and C_T the overall number of change-points observed until time T. We assume that the observations $x_t \sim \mathcal{B}(\mu_t)$ are generated by a piece-wise Bernoulli process such that there exists a non-decreasing change-points sequence $(\tau_c)_{c\in[1,C_T]} \in \mathbb{N}^{C_T}$ verifying:*

$$\begin{cases} \forall c \in [1, C_T]\,,\ \forall t \in \mathcal{T}_c = [\tau_c, \tau_{c+1}) \quad \mu_t = \theta_c, \\ \tau_1 = 1 < \tau_2 < \ldots < \tau_{C_T+1} = T+1. \end{cases} \tag{1}$$

Remark 1 (Interests in the Bernoulli case)**.** *The interests in working on the Bernoulli distributions are not as restrictive as it seems. On the first hand, from a concentration point of view, Bernoulli distributions can be seen as a worst case of bounded distributions. Moreover, Bernoulli distributions are crucially used in many widespread applications of machine learning. For instance:*

- *modelling the collisions in cognitive radio*
- *monitoring the performances of statistical models*
- *monitoring events in probes for network supervision*
- *the multi armed bandit problem*
- *experiments in clinical trials and recommender systems*

Notation 1. *In the following, we denote by $\mathbf{x}_{s:t} := (x_s, ..., x_t)$ the sequence of observations from time s up to time $t \geqslant s$. Furthermore, the length of the sequence $\mathbf{x}_{s:t}$ is denoted by $n_{s:t} := t - s + 1$ and the empirical mean over the sequence $\mathbf{x}_{s:t}$ is denoted by $\widehat{\mu}_{s:t} := \frac{1}{n_{s:t}} \sum_{i=s}^{t} x_i$.*

Definition 2 (Online change-point detection strategy)**.** *An online change-point strategy $\mathcal{A}$ takes sequentially as input a sequence $\mathbf{x}_{r:t}$ and output (at each time t) a binary scalar such that:*

$$\mathcal{A}(\mathbf{x}_{r:t}) = \begin{cases} 1 & \text{if a change in the generation of} \\ & \text{the sequence } \mathbf{x}_{r:t} \text{ is detected,} \\ 0 & \text{else.} \end{cases}$$

A strategy is said to be ***anytime*** *if it does not depend on the time horizon T which denotes the stream length.*

Performance assessment Let: $\mathbf{x}_{r:\tau_c-1} \sim \mathcal{B}(\theta_1)$, $\mathbf{x}_{\tau_c:t} \sim \mathcal{B}(\theta_2)$, τ_c the change-point to detect and r the starting time. The performance of an algorithm that aims at detecting the change-point $\tau_c \in [r, t]$ in the sequence $\mathbf{x}_{r:t}$ is assessed using two notions.

- **False alarm rate**: the probability of detecting a change at some instant $s \in [r, \tau_c)$ where there is no change. Usually, the false alarm rate is expressed as: $\mathbb{P}\Big\{\exists s \in [r, \tau_c) : \mathcal{A}(\mathbf{x}_{r:s}) = 1\Big\}$.
- **Detection delay**: the number of time steps needed to detect a change. It is formally defined for a strategy $\mathcal{A}$ as: $\widehat{\tau}_{\mathcal{A}}(\mathbf{x}_{r:t}) := \min\big\{s \in [r, t] : \mathcal{A}(\mathbf{x}_{r:s}) = 1\big\}$. Thus, the detection delay is expressed as: $\mathfrak{D}_{|\theta_2-\theta_1|,r,\tau_c} := (\widehat{\tau}_{\mathcal{A}}(\mathbf{x}_{r:t}) - \tau_c) \times \mathbb{I}\{\widehat{\tau}_{\mathcal{A}}(\mathbf{x}_{r:t}) > \tau_c\}$, where $\mathbb{I}\{\bullet\}$ denotes the indicator function.

Currently, the literature provides us with an interesting asymptotic lower bound on the **expected** detection delay. Theorem 1 gives the lower bound. (It is a reformulation of Theorem 3.1 in (Lai & Xing, 2010)).

Theorem 1 (Asymptotic lower bound on the expected detection delay)**.** *Let:* $\mathbf{x}_{r:\tau_c-1} \sim \mathcal{B}(\theta_1)$*,* $\mathbf{x}_{\tau_c:t} \sim \mathcal{B}(\theta_2)$*,* $\mathcal{A}$ *an online change-point detection strategy,* τ_c *the change-point to detect and* r *the starting time. Assuming that the false alarm rate is controlled such that:* $\mathbb{P}_{\theta_1}\Big\{\exists s \in [r, \tau_c) : \mathcal{A}(\mathbf{x}_{r:s}) = 1\Big\} \leqslant \delta$*, then as the quantity* $\frac{n_{r:\tau_c}}{|\log \delta|} \underset{\delta \to 0}{\to} \infty$*, the expected detection delay* $\mathbb{E}_{\theta_1,\theta_2}\left[\widehat{\tau}_{\mathcal{A}}(\mathbf{x}_{r:t}) - \tau_c\right]$ *is lower bounded as follows:*

$$\mathbb{E}_{\theta_1,\theta_2}\left[\widehat{\tau}_{\mathcal{A}}(\mathbf{x}_{r:t}) - \tau_c\right] \geqslant \left(\frac{\mathbb{P}_{\theta_1}\Big\{\widehat{\tau}_{\mathcal{A}}(\mathbf{x}_{r:t}) > \tau_c\Big\}}{\boldsymbol{kl}(\theta_2, \theta_1)}\right) \log \frac{1}{\delta}$$

where $\boldsymbol{kl}(\bullet, \bullet)$ *stands for the Kullback-Leibler divergence for Bernoulli distributions.*

3. The original Bayesian Online Change Point Detector (BOCPD)

In this section, we describe the original version of the Bayesian Online Change Point Detector introduced in (Fearnhead & Liu, 2007) and then revisited in (Adams & MacKay, 2007). Then,we reformulate it in term of a learning procedure using a growing number of forecaster. By this way, we highlight the difficulty of its analysis.

3.1. Learning using the runlength inference

(Adams & MacKay, 2007) have introduced an efficient Bayesian strategy for handling piece-wise stationary processes. This Bayesian strategy computes $p(r_t|\mathbf{x}_{1:t})$ the posterior distribution over the current runlength r_t, which denotes the number of time steps since the last change-point, given the data so far observed $\mathbf{x}_{1:t}$. The exact inference on the runlength distribution is done recursively using the following message-passing algorithm:

$$p(r_t|\mathbf{x}_{1:t}) \propto \sum_{r_{t-1}} \underbrace{p(r_t|r_{t-1})}_{\text{hazard}} \underbrace{p(x_t|r_{t-1}, \mathbf{x}_{1:t-1})}_{\text{UPM}} p(r_{t-1}|\mathbf{x}_{1:t-1})$$

where the hazard function is expressed as follows:

$$p(r_t|r_{t-1}) = \begin{cases} H(r_{t-1}) & \text{if } r_t = 0 \\ 1 - H(r_{t-1}) & \text{if } r_t = r_{t-1} + 1 \\ 0 & \text{otherwise} \end{cases} \quad (2)$$

with: $H(s) = \frac{\mathbf{P}_{\text{change}}(s+1)}{\sum_{t=s+1}^{\infty} \mathbf{P}_{\text{change}}(t)}$ and $\mathbf{P}_{\text{change}}$ denotes the probability distribution over the interval between changepoints.

A simple example of BOCPD would be to use a constant hazard function $h \in (0, 1)$ in the sense that $p(r_t = 0|r_{t-1})$ is independent of r_{t-1} and is constant, giving rise, a priori, to geometric inter-arrival times for change points ($\mathbf{P}_{\text{change}}(s + 1) = h(1 - h)^s$). Thus, the recursive runlength distribution computation becomes:

$$p(r_t \neq 0|\mathbf{x}_{1:t}) \propto (1 - h)\, p(x_t|r_{t-1}, \mathbf{x}_{1:t-1})p(r_{t-1}|\mathbf{x}_{1:t-1})$$
$$p(r_t = 0|\mathbf{x}_{1:t}) \propto h \sum_{r_{t-1}} p(x_t|r_{t-1}, \mathbf{x}_{1:t-1})p(r_{t-1}|\mathbf{x}_{1:t-1})$$

Then, for Bernoulli observations ($x_t \sim \mathcal{B}(\mu_t)$) the underlying predictive distribution (UPM) can be set to the Laplace predictor.

Definition 3 (Laplace predictor)**.** *The Laplace predictor* $\mathtt{Lp}(x_{t+1}|\mathbf{x}_{s:t})$ *takes as input a sequence* $\mathbf{x}_{s:t} \in \{0, 1\}^{n_{s:t}}$ *and predicts the value of the next observation* $x_{t+1} \in \{0, 1\}$ *as follows:*

$$\mathtt{Lp}(x_{t+1}|\mathbf{x}_{s:t}) := \begin{cases} \frac{\sum_{i=s}^{t} x_i + 1}{n_{s:t}+2} & \text{if } x_{t+1} = 1, \\ \frac{\sum_{i=s}^{t}(1-x_i)+1}{n_{s:t}+2} & \text{if } x_{t+1} = 0, \end{cases}$$

where $\forall x \in \{0, 1\}$ $\mathtt{Lp}(x|\emptyset) = \frac{1}{2}$ *corresponds to the uniform prior given to the process generating* θ_c*.*

Remark 2. *Laplace predictor is used as the estimator of the maximum likelihood with a uniform prior. It originates from the classical literature on universal codes and has standard robustness properties and Bayesian interpretation that make it of especial interest. Another variant is the Krichesky-Trofimov estimate (Cerqueira & Leonardi, 2018).*

Although BOCPD algorithm is very efficient in practice, its analysis in term of false alarm rate and detection delay is still an open problem. As a first step of the analysis of the Bayesian Online Change Point Detection, in this section we reformulate it in terms of learning strategy based on a growing number of forecasters.

In the following, to simplify the derivations (especially for Lemmas 1 and 2) we assume that the hazard function for BOCPD is constant ($H(r_{t-1}) = h$). Otherwise, the statement of Lemmas 1 and 2 becomes cumbersome to write and difficult to understand.

3.2. Learning with a growing number of forecasters

Notion of forecaster. Let $t \in \mathbb{N}^\star$ and $s \in [1, t]$. A forecaster s is a successive product of $(t - s)$ Laplace predictors $(\mathtt{Lp}(x_{t+1}|\mathbf{x}_{s:t}) \times \mathtt{Lp}(x_t|\mathbf{x}_{s:t-1}) \times ... \times \mathtt{Lp}(x_s|\emptyset))$ (see Definition 3), created at time s with some initial weight. At each time t, the forecaster s observes exactly the sequence $\mathbf{x}_{s:t}$ from the environment.

At each time t, each possible value of the runlength $r_t \in [0, t-1]$ corresponds to a specific forecaster. More specifically, the forecaster starting at time s corresponds at time t to the $t - s$ value of the runlength r_t.

Forecaster loss. Using the Laplace predictor, the instantaneous loss of the forecaster s at time t is given by:

$$\begin{aligned} l_{s,t} &:= -\log \mathtt{Lp}\left(x_t|\mathbf{x}_{s:t-1}\right) \\ &= -x_t \log \mathtt{Lp}\left(1|\mathbf{x}_{s:t-1}\right) - (1-x_t)\log \mathtt{Lp}\left(0|\mathbf{x}_{s:t-1}\right). \end{aligned}$$

Then, let $\widehat{L}_{s:t} := \sum_{s'=s}^{t} l_{s':t}$ denotes the cumulative loss incurred by the forecaster s from time s until time t which takes the following crude expression:

$$\widehat{L}_{s:t} := \sum_{s'=s}^{t} -\log \mathtt{Lp}\left(x_t|\mathbf{x}_{s':t-1}\right) \tag{3}$$

Forecaster weights. Instead of dealing with the posterior distribution of the runlength r_t, we propose to give to each forecaster s a weight $v_{s,t} := p(r_t = t-s|\mathbf{x}_{s:t})$ according to its sequence of observations $\mathbf{x}_{s:t}$. By this way, we describe the novel formulation of the Bayesian Online Change Point Detector in Algorithm 1. Notice that in line 5, Algorithm 1 performs a change-point detection, which was not present in (Adams & MacKay, 2007).

Algorithm 1 BOCPD (Fearnhead & Liu, 2007)

Input: $h \in (0,1)$

1: $v_{1,1} \leftarrow 1$
2: **for** $t = 1, \ldots$ **do**
3: Observe $x_t \sim \mathcal{B}\left(\mu_t\right)$
4: Define for each forecaster s up to time t:

$$v_{s,t} \leftarrow \begin{cases} (1-h)\exp\left(-l_{s,t}\right) v_{s,t-1} & \forall s < t, \\ h\sum_{i=1}^{t-1} \exp\left(-l_{i,t}\right) v_{i,t-1} & s = t\,. \end{cases} \tag{4}$$

5: Estimate the last change-point $\widehat{\tau}_t$ such that: $\widehat{\tau}_t \leftarrow \mathrm{argmax}_{s\in[1,t]}\, v_{s,t}$.
6: **end for**

Equation (4) defines the weights $v_{s,t}$ recursively. Lemma 1 expands the expression of $v_{s,t}$ for a better way to handle these quantities.

Lemma 1 (From recursive to closed-form expressions)**.** *Let: $V_t = \sum_{s=1}^{t} v_{s,t}$. Then, by noticing that $V_t = \sum_{s=1}^{t-1} \exp\left(-l_{s,t}\right) v_{s,t-1}$, the quantities $v_{s,t}$ take the following alternative closed-form expression:*

$$v_{s,t} = \begin{cases} (1-h)^{t-s+1}\, h^{\mathbb{I}\{s\neq 1\}} \exp\left(-\widehat{L}_{s:t}\right) V_s & \forall s < t, \\ hV_t & s = t. \end{cases}$$

First, from Lemma 1 one should notice that the quantity V_t plays the role of an initial weight that is given to the forecaster newly created at time t. Thus, in order to control the quantities $v_{t,s}$, we need to explicitly expand the expression of V_t.

The expression for V_t is given iteratively (see Lemma 1). Making it explicit reveals the power of the strategy introduced by (Fearnhead & Liu, 2007), that combines the updates of exponentially many forecasters into a simple iterative scheme. Indeed, Lemma 2 gives us the explicit expression of V_t.

Lemma 2 (Computing the initial weight V_t)**.** *The initial weight V_t takes the following form:*

$$V_t = (1-h)^{t-2} \sum_{k=1}^{t-1} \left(\frac{h}{1-h}\right)^{k-1} \tilde{V}_{k:t} \quad \textit{where:}$$

$$\tilde{V}_{k:t} = \sum_{i_1=1}^{t-k} \sum_{i_2=i_1+1}^{t-(k-1)} \ldots \sum_{i_{k-1}=i_{k-2}+1}^{t-2} \underline{\exp\left(-\widehat{L}_{1:i_1}\right)} \times$$
$$\prod_{j=1}^{k-2} \exp\left(-\widehat{L}_{i_j+1:i_{j+1}}\right) \exp\left(-\widehat{L}_{i_{k-1}+1:t-1}\right).$$

On the other hand, dealing with the explicit expression of V_t is challenging from a theoretical standpoint (proving performance guarantees), since we need to control a combinatorial number of cumulative losses. Indeed, notice that:

$$\sum_{i_1=1}^{t-k} \sum_{i_2=i_1+1}^{t-(k-1)} \ldots \sum_{i_{k-1}=i_{k-2}+1}^{t-2} 1 = \binom{t-2}{k-1},$$

where $\binom{\bullet}{\bullet}$ stands for the binomial operator.

4. The Restarted Bayesian Online Change Point Detector algorithm (R-BOCPD)

In this section, we introduce a pruning version of the original BOCPD which is built on a *novel initial weight function*, a *restart procedure to prune the useless experts* and a *well tuned hyper-parameter* instead of the hazard function.

4.1. Introducing a simple initial weight.

In order to avoid the difficulty mentioned in Lemma 2, we propose to use a much simpler initial weight that takes the following form:

$$\mathcal{V}_{r:s-1} := \exp\left(-\widehat{L}_{r:s-1}\right) \text{ for some starting time } r.$$

Notice that the initial weight $\mathcal{V}_{r:s-1}$ is a restricted version of the original one V_s by forgetting the contribution of all forecasters but the one launched at the starting time r (underlined term in Lemma 2). Thereby, the control of the initial weight is made easier: instead of dealing with a combinatorial number of cumulative losses, we only need to control **one** cumulative loss ($\widehat{L}_{r:s-1}$). Thus, for some starting time r, we denote by $\vartheta_{r,s,t}$ the novel weight given to the forecaster $s \geqslant r$ at time $t \geqslant s$. Then, one should also notice that BOCPD (Algorithm 1) produces an estimation of the last change point at each time step. In order to analyze this algorithm in term of detection delay, we propose to introduce a change-point decision rule (restart procedure).

4.2. Introducing a restart procedure.

For any starting time $r \leqslant t$, the change-point criterion is written as follows:

$$\texttt{\textbf{Restart}}_{r:t} = \mathbb{I}\{\exists s \in (r,t] : \vartheta_{r,s,t} > \vartheta_{r,r,t}\} \qquad (5)$$

where $\vartheta_{r,s,t}$ denotes the weight of the forecaster s created with the initial weight $\mathcal{V}_{r:t_{s-1}}$ at time t (see Algorithm 2 line 4). The intuition behind the criterion $\texttt{\textbf{Restart}}_{r:t}$ is that at each time $t < \tau$ where there is no change, the forecaster distribution tends to be concentrated around the forecaster launched at the starting time r. So, if the distribution $\vartheta_{r,s,t}$ undergoes a change then it can be seen as a certain change-point that has appeared. Thereby when $\texttt{\textbf{Restart}}_{r:t} = 1$, a change-point is detected and thus we **restart** a new forecaster at time $r = t+1$ and **delete** all previous launched forecasters $(s < r)$. This can be seen as a sophisticated pruning out procedure to reduce the number of launched forecasters.

Finally, by using the hyper-parameter $\eta_{r,s,t}$ (instead of the constant value h) and plugging the initial weight $\mathcal{V}_{r:t-1}$ and the decision rule $\texttt{\textbf{Restart}}_{r:t}$ into the formalism of BOCPD (Equation 4), we obtain a restarted version of the Bayesian Online Change Point Detector which is described in Algorithm 2.

Algorithm 2 R-BOCPD

Input: $\eta_{r,s,t} \in (0,1)$

1: $r \leftarrow 1$, $\vartheta_{r,1,1} \leftarrow 1$, $\eta_{r,1,1} \leftarrow 1$.
2: **for** $t = 1, \dots$ **do**
3: Observe $x_t \sim \mathcal{B}(\mu_t)$
4: Define for each forecaster s from time r to time t:

$$\vartheta_{r,s,t} \leftarrow \begin{cases} \frac{\eta_{r,s,t}}{\eta_{r,s,t-1}} \exp\left(-l_{s,t}\right) \vartheta_{r,s,t-1} & \forall s < t, \\ \eta_{r,t,t} \times \mathcal{V}_{r:t-1} & s = t\,. \end{cases} \qquad (6)$$

5: **if** $\texttt{\textbf{Restart}}_{r:t} = 1$ **then** $r \leftarrow t+1$, $\vartheta_{r,r,r} \leftarrow 1$, $\eta_{r,r,r} \leftarrow 1$.
6: Estimate the last change-point: $\widehat{\tau}_t \leftarrow r$.
7: **end for**

4.3. Discussion about R-BOCPD

The main difference between the R-BOCPD and its previous version lie primarily in the use of the test $\texttt{\textbf{Restart}}_{r:t} = 1$ for detecting the change-points. The second difference is the use of a simple initial weight $\mathcal{V}_{r:t-1}$ instead of the quantity V_t standing for the sum of the forecaster weights at time t. This is essentially done for theoretical reasons (see Lemma 2). The third difference is the use of a hyper-parameter $\eta_{r,s,t}$ instead of the hazard function h. The quantity $\frac{\eta_{r,s,t}}{\eta_{r,s,t-1}} \underset{t\to\infty}{\rightarrow} 1$ is used in updating the forecaster distribution $\vartheta_{r,s,t}$ at time t instead of the quantity $(1-h)$ whose asymptotic behavior is the same for an harazd rate very small. Indeed, $1-h \approx 1$. Finally, unlike (Adams & MacKay, 2007; Fearnhead & Liu, 2007) where the hazard function is assumed to be known (see Appendix A for more details), the function $\eta_{r,s,t}$ will be specified thanks to the analysis in section 5.1.

5. Performance guarantees

In this section, we build the two main guarantees for the R-BOCPD algorithm, namely: the false alarm rate (Theorem 2) and the detection delay control (Theorem 3). Then, we provide the reader with some useful tools to build these guarantees.

5.1. Non-asymptotic analysis of R-BOCPD

Let r denotes the time of the last restart and τ_c the change-point just coming after r.

Theorem 2 states the condition on $\eta_{r,s,t}$ where R-BOCPD (algorithm 2) does not make any false alarm with high probability. It is without reminding that the false alarm corresponds to the event where $\texttt{\textbf{Restart}}_{r:t} = 1$ during the period $[r, \tau_c)$.

Theorem 2 (False alarm rate)**.** *Assume that* $\mathbf{x}_{r:t} \sim \mathcal{B}(\theta)$. *Let:* $\alpha > 1$. *If* $\eta_{r,s,t}$ *is small enough such that:*

$$\forall t \in [r, \tau_c), s \in (r,t] :$$
$$\eta_{r,s,t} < \frac{\sqrt{n_{r:s-1} \times n_{s:t}}}{10\,(n_{r:t}+1)} \times \left(\frac{\log(4\alpha+2)\delta^2}{4n_{r:t}\log((\alpha+3)\,n_{r:t})}\right)^{\alpha}$$

then, with probability higher than $1-\delta$, *no false alarm occurs on the interval* $[r, \tau_c)$*:*

$$\mathbb{P}_\theta\Big\{\exists\, t \in [r, \tau_c) : \texttt{\textbf{Restart}}_{r:t} = 1\Big\} \leqslant \delta.$$

Before stating the control of the detection delay, we need to introduce the notion of relative gap $\Delta_{r,s,t}$.

Definition 4 (Relative gap $\Delta_{r,s,t}$)**.** *Let* $\Delta \in [0,1]$. *The relative gap* $\Delta_{r,s,t}$ *for the forecaster* s *at time* t *takes the following form (depending on the position of* s*):*

$$\Delta_{r,s,t} = \left(\frac{n_{r:\tau_c-1}}{n_{r:s-1}}\mathbb{I}\{\tau_c \leqslant s \leqslant t\} + \frac{n_{\tau_c:t}}{n_{s:t}}\mathbb{I}\{s < \tau_c\}\right)\Delta.$$

Theorem 3 states the detection delay under some condition on the quantity $\eta_{r,s,t}$.

Theorem 3 (Detection delay)**.** *Let* $\mathbf{x}_{r:\tau_c-1} \sim \mathcal{B}(\theta_1)$, $\mathbf{x}_{\tau_c:t} \sim \mathcal{B}(\theta_2)$ *and* $\Delta = |\theta_1 - \theta_2|$*: the change-point gap. Then, let:* $f_{r,s,t} = \log n_{r:s} + \log n_{s:t+1} - \frac{1}{2}\log n_{r:t} + \frac{9}{8}$. *If* $\eta_{r,s,t}$ *is large enough such that:*

$$\eta_{r,s,t} > \exp\big(-2n_{r,s-1}\left(\Delta_{r,s,t} - \mathcal{C}_{r,s,t,\delta}\right)^2 + f_{r,s,t}\big),$$

then, the change-point τ_c *is detected (with a probability at* *least* $1-\delta$*) with a delay not exceeding* $\mathfrak{D}_{\Delta,r,\tau_c}$*, such that:*

$$\mathfrak{D}_{\Delta,r,\tau_c} = \min\left\{d \in \mathbb{N}^\star : d > \frac{\left(1-\frac{\mathcal{C}_{r,\tau_c,d+\tau_c-1,\delta}}{\Delta}\right)^{-2}}{2\Delta^2}\times \frac{-\log\eta_{r,\tau_c,d+\tau_c-1}+f_{r,\tau_c,d+\tau_c-1}}{1+\frac{\log\eta_{r,\tau_c,d+\tau_c-1}-f_{r,\tau_c,d+\tau_c-1}}{2n_{r,\tau_c-1}(\Delta-\mathcal{C}_{r,\tau_c,d+\tau_c-1,\delta})^2}}\right\}, \quad (7)$$

where: $$\mathcal{C}_{r,s,t,\delta} = \frac{\sqrt{2}}{2}\left(\sqrt{\frac{1+\frac{1}{n_{r:s-1}}}{n_{r:s-1}}\log\left(\frac{2\sqrt{n_{r:s}}}{\delta}\right)} + \sqrt{\frac{1+\frac{1}{n_{s:t}}}{n_{s:t}}\log\left(\frac{2n_{r:t}\sqrt{n_{s:t}+1}\log^2(n_{r:t})}{\log(2)\delta}\right)}\right). \quad (8)$$

Discussion about the detection delay $\mathfrak{D}_{\Delta,r,\tau_c}$ From Eq. (7), for a fixed r and τ_c we notice that the larger the change-point gap Δ, the smaller the detection delay $\mathfrak{D}_{\Delta,r,\tau_c}$ and vice versa. Moreover for a fixed gap Δ, the larger n_{r,τ_c-1}: the number of observations before the change-point τ_c, the smaller the detection delay $\mathfrak{D}_{\Delta,r,\tau_c}$ (cf figure 1).

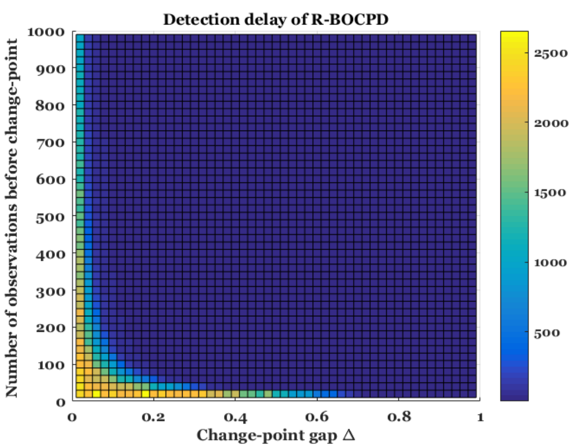

Figure 1. Variation of the R-BOCPD detection delay $\mathfrak{D}_{\Delta,r,\tau_c}$ as a function of the change point gap Δ (x-axis) and the number of observations before the change-point n_{r,τ_c-1} (y-axis). For this plot, we choose $\eta_{r,s,t}=\frac{1}{n_{r:t}}$ for R-BOCPD.

Remark 3 (Minimum detectable gap $\Delta_{\text{min}}(r,\tau_c,t)$)**.** *Instead of imposing a condition on the lower-bound of* $\eta_{r,s,t}$*, we can discuss the detectability of the change-point* τ_c *according to the magnitude of the gap* Δ*. Thus, if the gap* Δ *is of magnitude at least* $\Delta_{min}(r,\tau_c,t) = \sqrt{\frac{-\log\eta_{r,\tau_c,t}+f_{r,\tau_c,t}}{n_{r:\tau_c-1}}}+\mathcal{C}_{r,\tau_c,t,\delta}$*, then the change-point* τ_c *is detected (with a probability at least* $1-\delta$*) with a finite delay not exceeding* $\mathfrak{D}_{\Delta,r,\tau_c}$*.*

Discussion about the asymptotic optimality We compare the result of Theorem 3 with the existing lower-bound on the detection delay (see (Lai & Xing, 2010)). The asymptotic regime corresponds to the case where the elapsed time between the last restart r and the new change point τ_c tends to infinity, while the probability of false alarm δ tends to zero. More formally, the asymptotic regime is reached when $\frac{n_{r:\tau_c-1}}{\log(1/\delta)}\to\infty$, and $\log n_{r:\tau_c-1} = o\big(\log\frac{1}{\delta}\big)$ when $\delta\to 0$. We obtain that:

$$\mathfrak{D}_{\Delta,r,\tau_c} \underset{\tau_c\to\infty}{\to} \frac{-\log\eta_{r,\tau_c,d+\tau_c-1}+o\big(\log\frac{1}{\delta}\big)}{2\left|\theta_2-\theta_1\right|^2}. \quad (9)$$

Thus, following Theorem 1, the detection delay $\mathfrak{D}_{\Delta,r,\tau_c}$ is ***asymptotically order optimal*** (up to the Pinsker inequality tightness relating $|\theta_2-\theta_1|^2$ to $\boldsymbol{kl}(\theta_2,\theta_1)$). Moreover, the smaller $\eta_{r,s,t}$, the larger the detection delay and vice-versa. Also, following Remark 3, the smaller $\eta_{r,s,t}$, the larger the minimum detectable gap $\Delta_{\text{min}}(r,\tau_c,t)$.

Remark 4 (Main difficulty to get optimality with the Kullback-Leibler divergence)**.** *The result of Equation 9 shows the Euclidean distance* $|\theta_2-\theta_1|^2$ *instead of the Kullback-Leibler divergence* $\boldsymbol{kl}(\theta_2,\theta_1)$ *as expected from Theorem 1. Indeed, in the analysis of the detection delay (see Appendix C), the quantity* $n_{r:s-1}\boldsymbol{kl}(\widehat{\mu}_{r:s-1},\widehat{\mu}_{r:t})+n_{s:t}\boldsymbol{kl}(\widehat{\mu}_{s:t},\widehat{\mu}_{r:t})$ *appears (it naturally comes from Lemma 3). Building a lower bound of* $n_{r:s-1}\boldsymbol{kl}(\widehat{\mu}_{r:s-1},\widehat{\mu}_{r:t})+n_{s:t}\boldsymbol{kl}(\widehat{\mu}_{s:t},\widehat{\mu}_{r:t})$ *showing the quantity* $\boldsymbol{kl}(\theta_2,\theta_1)$ *(in the case where there is change point* $\tau\in[r,t]$*) does not seem trivial. Thus, we have opted to use the Pinsker inequality which has slightly reduced the optimality of our result but it remains significant.*

Discussion about the choice of $\eta_{r,s,t}$**.** Choosing $\eta_{r,s,t}\approx\frac{1}{n_{r:t}}$ seems to be a wise choice since it allows us to verify the conditions of Theorem 2 (for some $\alpha\to 1$) and Theorem 3. Thus, by plugging $\eta_{r,s,t}\approx\frac{1}{n_{r:t}}$ into the asymptotic expression of the detection delay (Equation 9), we get (in the asymptotic regime $\frac{n_{r:\tau_c-1}}{\log(1/\delta)}\to\infty$, and $\log n_{r:\tau_c-1} = o\big(\log\frac{1}{\delta}\big)$ when $\delta\to 0$):

$$\mathfrak{D}_{|\theta_2-\theta_1|,r,\tau_c} \underset{\tau_c\to\infty}{\to} \frac{o\big(\log\frac{1}{\delta}\big)}{2\left|\theta_2-\theta_1\right|^2}.$$

By this way, the detection delay is **asymptotically optimal** in the sense of Theorem 1 (up to the Pinsker inequality tightness relating $|\theta_2-\theta_1|^2$ to $\boldsymbol{kl}(\theta_2,\theta_1)$).

5.2. Sketch of proof for the false alarm rate control and the detection delay

In this section, we provide the key elements and the essential intuitions to build the false alarm guarantee and the detection delay control. The key element in building the false alarm guarantee and the detection delay of R-BOCPD lies in controlling *efficiently* the quantities $\log\vartheta_{r,s,t}$. Indeed using Equation (6), we get (for some starting time r):

$$\log\vartheta_{r,s,t} = \log\eta_{r,s,t}\times\mathbb{I}\{s\neq r\} - \widehat{L}_{r:s-1} - \widehat{L}_{s:t}. \quad (10)$$

Then, it is clear that controlling the quantity $\log \vartheta_{r,s,t}$ requires an efficient control of $\widehat{L}_{s:t}$.
Using the crude expression of $\widehat{L}_{s:t}$ (see Equation (3)) seems to be very naive in the sense that we need to control each instantaneous loss $l_{s,t}$ independently without taking into account the dependencies between $l_{s,t}$ and $l_{s,t-1}$. A smarter way to deal with the quantity $\widehat{L}_{s:t}$ lies in writing it as follows:

Lemma 3 (Rewriting the cumulative loss)**.** *Based on the Laplace predictor, the cumulative loss $\widehat{L}_{s:t}$ takes the following closed-form expression:*

$$\forall \mathbf{x}_{s:t} \in \{0,1\}^{n_{s:t}}\ \widehat{L}_{s:t} = \log\left(n_{s:t}+1\right) + \log\begin{pmatrix} n_{s:t} \\ \sum_{i=s}^{t} x_i \end{pmatrix}.$$

Remark 5. *The main idea of Lemma 3 is taken from the book "Prediction, Learning and Games" by (Cesa-Bianchi & Lugosi, 2006). An original proof by induction is provided in the appendix. Notice that Lemma 3 is valid for any binary sequence $\mathbf{x}_{s:t} \in \{0,1\}^{n_{s:t}}$. No assumption on the intrinsic distribution of the sequence $\mathbf{x}_{s:t}$ is required.*

Lemma 4 (Cumulative loss control before a change-point)**.** *Assume that we observe a sequence $\mathbf{x}_{s:t} \sim \mathcal{B}(\theta)$. Then, the control of the quantity $\widehat{L}_{s:t}$ is done as follows:*

Upper bound:

$$\widehat{L}_{s:t} \leqslant \log n_{s:t+1} - \sum_{i=s}^{t} x_i \log\theta - \sum_{i=s}^{t}(1-x_i)\log(1-\theta),$$

Lower bound:

$$\widehat{L}_{s:t} \geqslant -\sum_{i=s}^{t} x_i \log\theta - \sum_{i=s}^{t}(1-x_i)\log(1-\theta) + \log\frac{n_{s:t+1}}{\sqrt{n_{s:t}}} - n_{s:t}\boldsymbol{kl}\left(\widehat{\mu}_{s:t},\theta\right) - \frac{9}{8}.$$

Remark 6. *Notice how tight are the upper and lower bound of the loss $\widehat{L}_{s:t}$. The control in Lemma 4 represents an essential element to provide the false alarm guarantee in Theorem 2.*

Finally, one should notice that the lower bound of the cumulative loss $\widehat{L}_{s:t}$ involves the Kullback-Leibler divergence $\boldsymbol{kl}\left(\widehat{\mu}_{s:t},\theta\right)$. For very tight control of the cumulative loss, we need to efficiently control the quantity $\boldsymbol{kl}\left(\widehat{\mu}_{s:t},\theta\right)$. This is done uniformly in Lemma 5 and Lemma 6.

Lemma 5 (Time uniform $\boldsymbol{kl}(\bullet,\bullet)$ concentration)**.** *Let: $\widehat{\mu}_t$ denotes the empirical mean over the sequence $x_1,\dots,x_t \sim \mathcal{B}(\theta)$, then for all $(\delta,\alpha) \in (0,1)\times(1,\infty)$ we have:*

$$\mathbb{P}_\theta\left\{\forall t \in \mathbb{N}^\star : \boldsymbol{kl}\left(\widehat{\mu}_t,\theta\right) < \frac{\alpha}{t}\log\frac{\log(\alpha t)\log(t)}{\log^2(\alpha)\delta}\right\} \geqslant 1-\delta.$$

Lemma 6 (Doubly-time uniform $\boldsymbol{kl}(\bullet,\bullet)$ concentration)**.** *Let: $\widehat{\mu}_{s:t}$ denotes the empirical mean over the sequence $x_s,\dots,x_t \sim \mathcal{B}(\theta)$, then for all $(\delta,\alpha) \in (0,1)\times(1,\infty)$ we have:*

$$\mathbb{P}_\theta\Big\{\forall t \in \mathbb{N}^\star, \forall s \in (r,t] : \boldsymbol{kl}\left(\widehat{\mu}_{s:t},\theta\right) < \frac{\alpha}{n_{s:t}} \times \log\frac{n_{r:t}\log^2(n_{r:t})\log((\alpha+1)\,n_{s:t})}{\log(2)\log^2(\alpha)\delta}\Big\} \geqslant 1-\delta.$$

Then, in order to build the detection delay guarantee, we will need to efficiently control the quantity $|\widehat{\mu}_{r:s-1} - \widehat{\mu}_{s:t}|$ (which is related to $\widehat{L}_{r:s-1} + \widehat{L}_{s:t} - \widehat{L}_{r:t}$ via Pinsker inequality). This is done thanks to Lemma 7.

Lemma 7 (Doubly-time uniform concentration)**.** *Let: $x_r, \dots x_t$ be a sequence of independent binary random variables sampled from a Bernoulli distribution and $\widehat{\mu}_{i:j}$ the empirical mean over the sequence $\mathbf{x}_{i:j}$. Then, for all $(r,\delta) \in \mathbb{N}^\star \times (0,1)$, we get the following control:*

$$\mathbb{P}\Big\{\exists\, t > r, s \in [r,t) : |\widehat{\mu}_{r:s-1} - \widehat{\mu}_{s:t} - \mathbb{E}\left[\widehat{\mu}_{r:s-1} - \widehat{\mu}_{s:t}\right]| \geqslant \mathcal{C}_{r,s,t,\delta}\Big\} \leqslant \delta,$$

where: $\mathcal{C}_{r,s,t,\delta}$ is defined in Equation (8).

6. Numerical illustrations of R-BOCPD against the state-of-art

In this section, we provide numerical comparisons between the proposed strategy R-BOCPD and two state-of-art strategies: BOCPD and the Improved GLR (Maillard, 2019) in two different schemes, a first comparison on synthetic data and a second comparison on real world data. Software and simulation code is available at `https://github.com/Ralami1859/Restarted-BOCPD`.

6.1. Synthetic data

6.1.1. Comparison with BOCPD

Highlighting the use of the function $\mathcal{V}_{r:t-1}$ In order to highlight the use of the function $\mathcal{V}_{r:t-1}$ as initial weight given to the forecaster newly created at time t (instead of the original one V_t), we compare the R-BOCPD strategy against its previous version BOCPD in the following experimental setting. We generate 2500 trajectories (sequences) of length $T = 5000$ where we vary the number of observation before the change-point from 10 to 1000 and we vary the change-point gap Δ from 0.01 to 1.

Then, we run R-BOCPD and BOCPD strategy on the same sequence 600 times. Finally, we plot the mean detection delay difference between R-BOCPD and BOCPD. Each square corresponds to a detection event for a change-point τ_c. The y coordinate corresponds to the number of observations both R-BOCPD and BOCPD algorithms received before the change-point, the x coordinate is the gap of the change-point. From figure 3(a), we see that the detection delay of R-BOCPD is slightly smaller than the one of BOCPD. Indeed, the detection delay difference is negative over all the experiments. By the way, using the function $\mathcal{V}_{r:t-1}$ instead of V_t as an initial weight given to the forecaster newly created allows us to speed up the change-point detection.

Highlighting the use of the restart procedure $\texttt{Restart}_{r:t}$ In order to highlight the benefit of using the restart procedure $\texttt{Restart}_{r:t}$ in R-BOCPD, we compare R-BOCPD strategy against BOCPD in the following setting. We generate a piece-wise stationary Bernoulli process with four change-points observed at time $(\tau_1 = 1, \tau_2 = 301, \tau_3 = 701, \tau_4 = 1051)$, then we run R-BOCPD and BOCPD on this environment and finally we plot (in figure 2) the change-point estimation $\widehat{\tau}_t$ for both R-BOCPD and BOCPD.

From figure 2, the curves of R-BOCPD and BOCPD are almost the same meaning that there is no significant difference in terms of false alarm and detection delays for both algorithms. Thus, restart procedure $\texttt{Restart}_{r:t}$ doesn't affect the performances of the Bayesian online change-point detector.

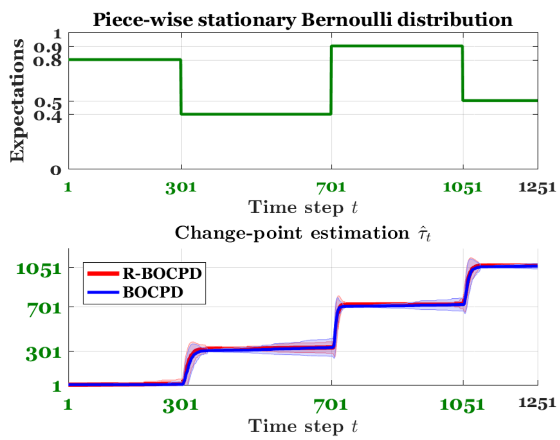

Figure 2. In all the experiment, we choose $\eta_{r,s,t} = \frac{1}{n_{r:t}}$ for R-BOCPD and $h = 3/1200$ for BOCPD. The curves are averaged over 300 runs. (Their error bars are also plotted).

6.1.2. Comparison with the Improved GLR

Recently, the classical Generalized Likelihood Ratio (GLR) strategy has been improved by (Maillard, 2019) by well-tuning the decision threshold. It used a novel sharp concentration inequality based on the Laplace method for scan-statistics which holds doubly uniformly in time (see Lemma 7). The final formulation of the Improved GLR strategy for Bernoulli processes takes the following form:

$$\mathrm{GLR}_{r:t} = \mathbb{I}\left\{\exists\, s \in [r, t) : \left|\widehat{\mu}_{r:s} - \widehat{\mu}_{s+1:t}\right| \geqslant \mathcal{C}_{r,s,t,\delta}\right\}$$

where $\mathcal{C}_{r,s,t,\delta}$ is defined in Equation (8).

This strategy has been proven to be asymptotically order optimal, in the sense of Theorem 1 (see Theorem 6 in (Maillard, 2019)). Therefore, comparing R-BOCPD against the Improved GLR strategy is a wise choice since GLR is considered as a very good baseline for the setting of the paper. Thus in Figure 3(b), we generate 2500 trajectories (sequences) of length $T = 2500$ where we vary the number of observation before the change-point from 10 to 500 and we vary the change-point gap Δ from 0.01 to 1. Then, we run R-BOCPD and Improved GLR strategy on the same sequence 360 times. Finally, we plot the mean detection delay difference between R-BOCPD and Improved GLR.

Figure 3(b) highlights the benefit of the R-BOCPD algorithm over the Improved GLR strategy. Indeed, the detection delay of R-BOCPD is slightly smaller than the one of the Improved GLR strategy. The while square means that Improved GLR isn't able to perform a detection while R-BOCPD does. Thus, for the small gap case, R-BOCPD is more robust than the Improved GLR strategy.

6.2. Experiments on Well-log data (Real world data)

These data have been studied in the context of change-point detection by (Fearnhead & Clifford, 2003) and has become a benchmark data set for uni-variate change-point detection. They consist on 4050 measurements $\left(y_1, \ldots y_{4050} \in \left[6.42 \times 10^4, 1.04 \times 10^5\right]\right)$ of nuclear magnetic response taken during the drilling of a well. The data are used to interpret the geophysical structure of the rock surrounding the well. The variations in mean reflect the stratification of the earth's crust. In order to perform the experiments on the Well-log data, we typically proceed by computing the re-scaled values $\tilde{y}_1, \ldots, \tilde{y}_{4050} \in [0, 1]$ then we use the sequence of samples $x_1 \sim \mathcal{B}\left(\tilde{y}_1\right) \ldots, x_{4050} \sim \mathcal{B}\left(\tilde{y}_{4050}\right)$ as input for BOCPD, R-BOCPD and GLR (see figure 4). Finally, we plot the estimation of the change-points for each algorithm. Notice that R-BOCPD is more suitable to detect the change-point at $t = 2800$ than BOCPD or GLR.

7. Extension to finite support distributions

In this paper, the restarted Bayesian Online Change-point detector (R-BOCPD) has been specially designed for Bernoulli distributions (see Lemma 3 and Definition 3). This does not seem to have much limited its applications. Extensions to discrete-support distributions might be possible, resorting for instance to other concentration inequalities results, but this is not the purpose of this work. Note that, algorithms working on Bernoulli distributions can be conveniently extended to work with other set of distributions with bounded support: A classical way to do so when considering distribution $\mathcal{D}$ with bounded support in $[a, b]$, and some observation $y_t \sim \mathcal{D} \in [a, b]$, is to compute the re-scaled observation $\tilde{y}_t \in [0, 1]$, then use the sample $x_t \sim \mathcal{B}\left(\tilde{y}_t\right)$ as input (This transformation may not preserve optimality

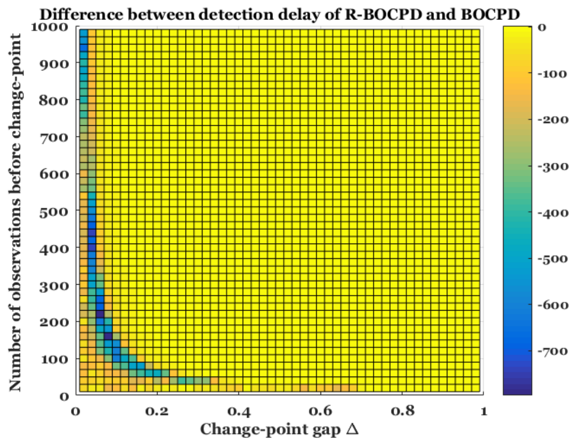

(a) Difference between detection delays of R-BOCPD and BOCPD.

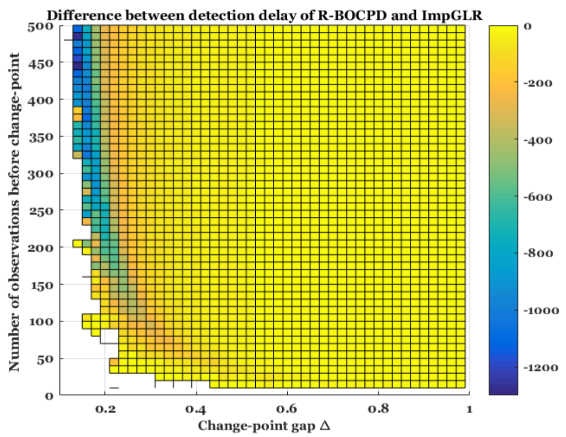

(b) Difference between detection delays of R-BOCPD and GLR. The white square means that ImprGLR isn't able to perform a detection while R-BOCPD does.

Figure 3. In all the experiments, we choose $\eta_{r,s,t} = \frac{1}{n_{r:t}}$ for R-BOCPD and $h = 1/T$ for BOCPD. The parameter δ (false alarm rate of ImpGLR) is set to 0.01.

properties, though). The experiments in section 6.2 has been performed following this procedure.

It should be noted that our analysis makes use of some specific properties of Bernoulli distributions, such as concentration inequalities, the key Lemma 3 and existence of explicit conjugate priors. Lemmas 5, 6 and 7 on the other hand are valid for any sub-Gaussian distributions. The Extension of our analysis to other popular distributions (e.g. Gaussians, Poisson, etc.) would need the specific equivalent of Lemma 3, and depend on the specific concentration inequalities and conjugate priors for this family. We believe this requires careful case by case examination.

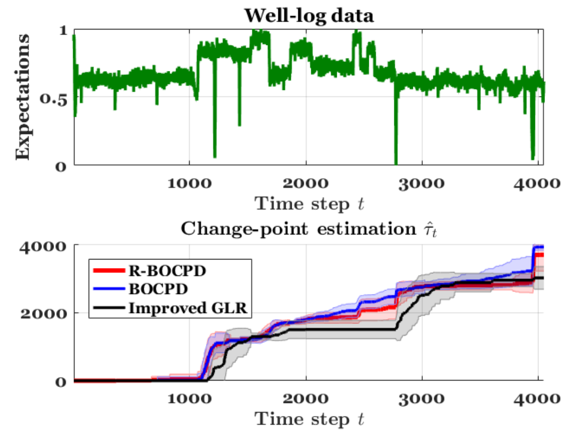

Figure 4. In all the experiment, we choose $\eta_{r,s,t} = \frac{1}{n_{r:t}}$ for R-BOCPD and $h = 1/T$ for BOCPD. The curves are averaged over 300 runs. (Their error bars are also plotted).

8. Conclusion and future work

In this paper, we introduced an improvement of the Bayesian Online Change-point Detector, called Restarted BOCPD. We provided a non-asymptotic analysis of its false alarm rate and detection delay, and shown numerically that our proposal outperforms its previous version. This constitutes arguably the first problem-dependent analysis of a Bayesian strategy in the context of change-point detection, and opens the path towards a complete analysis of BOCPD on the one hand, and the development of other Bayesian alternative on the other hand. We note that obtaining such guarantees is of primary importance in some applications, and in particular in the increasingly popular context of non-stationary multi-armed bandits.

Acknowledgments

This work has been supported by the French Agence Nationale de la Recherche (ANR), under grant ANR-16- CE40-0002 (project BADASS).

References

Adams, R. P. and MacKay, D. J. Bayesian online changepoint detection. *arXiv preprint arXiv:0710.3742*, 2007.

Alami, R., Maillard, O., and Féraud, R. Memory bandits: a bayesian approach for the switching bandit problem. In *Neural Information Processing Systems: Bayesian Optimization Workshop.*, 2016.

Aminikhanghahi, S. and Cook, D. J. A survey of methods for time series change point detection. *Knowledge and Information Systems*, 51(2):339–367, May 2017. ISSN 0219-3116. doi: 10.1007/s10115-016-0987-z. URL https://doi.org/10.1007/s10115-016-0987-z.

Basseville, M., Nikiforov, I. V., et al. *Detection of abrupt changes: theory and application*, volume 104. Prentice Hall Englewood Cliffs, 1993.

Biswas, R., Limketkai, B., Sanner, S., and Thrun, S. Towards object mapping in non-stationary environments with mobile robots. In *IEEE/RSJ international conference on intelligent robots and systems*, volume 1, pp. 1014–1019. IEEE, 2002.

Brodsky, E. and Darkhovsky, B. S. *Nonparametric Methods in Change-Point Problems*, volume 243. Springer, 1993.

Caron, F., Doucet, A., and Gottardo, R. On-line changepoint detection and parameter estimation with application to genomic data. *Statistics and Computing*, 22(2):579–595, 2012.

Cerqueira, A. and Leonardi, F. Strong consistency of krichevsky-trofimov estimator for the number of communities in the stochastic block model. *arXiv preprint arXiv:1804.03509*, 2018.

Cesa-Bianchi, N. and Lugosi, G. *Prediction, learning, and games*. Cambridge university press, 2006.

Csörgö, M. and Horváth, L. *Limit theorems in change-point analysis*, volume 18. John Wiley & Sons Inc, 1997.

Dakdouk, H., Tarazona, E., Alami, R., Féraud, R., Papadopoulos, G. Z., and Maillé, P. Reinforcement learning techniques for optimized channel hopping in ieee 802.15.4-tsch networks. In *Proceedings of the 21st ACM International Conference on Modeling, Analysis and Simulation of Wireless and Mobile Systems*, MSWIM '18, pp. 99–107, New York, NY, USA, 2018. ACM. ISBN 978-1-4503-5960-3. doi: 10.1145/3242102.3242110.

Fearnhead, P. and Clifford, P. On-line inference for hidden markov models via particle filters. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 65 (4):887–899, 2003.

Fearnhead, P. and Liu, Z. On-line inference for multiple changepoint problems. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 69(4):589–605, 2007.

Goldberg, D. and Matarić, M. J. Maximizing reward in a non-stationary mobile robot environment. *Autonomous Agents and Multi-Agent Systems*, 6(3):287–316, 2003.

Grzegorczyk, M. and Husmeier, D. Non-stationary continuous dynamic bayesian networks. In *Advances in Neural Information Processing Systems*, pp. 682–690, 2009.

Jie, C. and Gupta, A. Parametric statistical change point analysis. *Birkh User*, 2000.

Kerkouche, R., Alami, R., Féraud, R., Varsier, N., and Maillé, P. Node-based optimization of lora transmissions with multi-armed bandit algorithms. In *2018 25th International Conference on Telecommunications (ICT)*, pp. 521–526, June 2018. doi: 10.1109/ICT.2018.8464949.

Knoblauch, J. and Damoulas, T. Spatio-temporal bayesian on-line changepoint detection with model selection. *arXiv preprint arXiv:1805.05383*, 2018.

Knoblauch, J., Jewson, J. E., and Damoulas, T. Doubly robust bayesian inference for non-stationary streaming data with $beta$-divergences. In *Advances in Neural Information Processing Systems*, pp. 64–75, 2018.

Konidaris, G., Kuindersma, S., Grupen, R., and Barto, A. G. Constructing skill trees for reinforcement learning agents from demonstration trajectories. In *Advances in neural information processing systems*, pp. 1162–1170, 2010.

Lai, T. L. and Xing, H. Sequential change-point detection when the pre-and post-change parameters are unknown. *Sequential analysis*, 29(2):162–175, 2010.

Maillard, O.-A. Sequential change-point detection: Laplace concentration of scan statistics and non-asymptotic delay bounds. In *Algorithmic Learning Theory*, pp. 610–632, 2019.

Mellor, J. and Shapiro, J. Thompson sampling in switching environments with bayesian online change detection. In *Artificial Intelligence and Statistics*, pp. 442–450, 2013.

Nandhini, V. and Devasena, M. G. Predictive analytics for climate change detection and disease diagnosis. In *2019 5th International Conference on Advanced Computing & Communication Systems (ICACCS)*, pp. 1152–1155. IEEE, 2019.

Niekum, S., Osentoski, S., Atkeson, C. G., and Barto, A. G. Champ: Changepoint detection using approximate model parameters. Technical report, CARNEGIE-MELLON UNIV PITTSBURGH PA ROBOTICS INST, 2014.

Page, E. S. Continuous inspection schemes. *Biometrika*, 41 (1/2):100–115, 1954.

Panda, S. P. and Nayak, A. K. Automatic speech segmentation in syllable centric speech recognition system. *International Journal of Speech Technology*, 19(1):9–18, 2016.

Polunchenko, A. S., Tartakovsky, A. G., and Mukhopadhyay, N. Nearly optimal change-point detection with an application to cybersecurity. *Sequential Analysis*, 31(3): 409–435, 2012.

Ruggieri, E. and Antonellis, M. An exact approach to bayesian sequential change point detection. *Computational Statistics & Data Analysis*, 97:71–86, 2016.

Saatçi, Y., Turner, R. D., and Rasmussen, C. E. Gaussian process change point models. In *ICML*, pp. 927–934, 2010.

Schmitt, T. A., Chetalova, D., Schäfer, R., and Guhr, T. Non-stationarity in financial time series: Generic features and tail behavior. *EPL (Europhysics Letters)*, 103(5):58003, 2013.

Tartakovsky, A. Sequential methods in the theory of information systems, 1991.

Turner, R., Saatci, Y., and Rasmussen, C. E. Adaptive sequential bayesian change point detection. In *Temporal Segmentation Workshop at NIPS*, pp. 1–4, 2009.

Turner, R. D., Bottone, S., and Stanek, C. J. Online variational approximations to non-exponential family change point models: with application to radar tracking. In *Advances in Neural Information Processing Systems*, pp. 306–314, 2013.

Wilson, R. C., Nassar, M. R., and Gold, J. I. Bayesian online learning of the hazard rate in change-point problems. *Neural computation*, 22(9):2452–2476, 2010.

Wu, Y. *Inference for change point and post change means after a CUSUM test*, volume 180. Springer, Lecture Notes in Statistics, 2007.

Xuan, X. and Murphy, K. Modeling changing dependency structure in multivariate time series. In *Proceedings of the 24th international conference on Machine learning*, pp. 1055–1062, 2007.