Ensemble binary segmentation for irregularly spaced data with change-points

Karolos K. Korkas¹

Received: 17 July 2020 / Accepted: 1 April 2021 / Published online: 15 April 2021
© The Author(s) 2021

Abstract

We propose a new technique for consistent estimation of the number and locations of the change-points in the structure of an irregularly spaced time series. The core of the segmentation procedure is the ensemble binary segmentation method (EBS), a technique in which a large number of multiple change-point detection tasks using the binary segmentation method are applied on sub-samples of the data of differing lengths, and then the results are combined to create an overall answer. We do not restrict the total number of change-points a time series can have, therefore, our proposed method works well when the spacings between change-points are short. Our main change-point detection statistic is the time-varying autoregressive conditional duration model on which we apply a transformation process in order to decorrelate it. To examine the performance of EBS we provide a simulation study for various types of scenarios. A proof of consistency is also provided. Our methodology is implemented in the R package `eNchange`, available to download from CRAN.

Keywords Multiple change-point detection · Conditional duration models · Binary segmentation · Ensemble methods · Non-stationarity

1 Introduction

Irregularly spaced time series, i.e. data recorded as and when they emerge, have attracted significant attention due to the development and increasing automation in information technology. This type of data, aka high-frequency data, has not only penetrated financial markets, whereby trade transactions, as a result of continuous electronic trading, have generated a vast amount of datasets. Examples of irregularly spaced time series can be found in many industrial and scientific domains such as in natural disasters where floods, earthquakes or volcanic eruptions typically occur at irregular time intervals; in clinical trials where a patient's state of health is typically

✉ Karolos K. Korkas
kkorkas@yahoo.co.uk

¹ School of Computing and Mathematical Sciences, University of Greenwich, London, UK

observed at different points of time; or in data collecting sensors (smart applications of Internet of Things) which are activated only when an event happens to, e.g., preserve battery energy.

Modelling high-frequency data successfully requires that the dynamics in the time series data are captured efficiently and standard time series methods are not appropriate for these data. To address this, a class of models, the autoregressive conditional duration models (henceforth, ACD), was originally introduced by Engle and Russell (1998). In particular, the authors model the conditional mean as an autoregressive process which has a multiplicative relationship with positive valued error terms. These specifications resemble the multiplicative structure of the GARCH model used for modelling an asset's return volatility, but thanks to the positive valued error terms ACDs are more suitable for modelling the dynamics of continuous positive valued random variables, such as durations (the time it takes between two trades occurring in a trading book), trading volume (the size of orders), or market depth (the available liquidity at any point of time). ACD models have become popular with many variations owing partially to the popularity of the GARCH model and its numerous extensions. An alternative to ACD's autoregressive specification is to parameterize the intensity function which assumes a self-exciting process. This class of stochastic process was originated by Hawkes (1971) and, hence, these stochastic processes are referred to as Hawkes processes. They serve as epidemic models: the occurrence of a number of events (e.g. seismic or buy/sell) will increase the probability of further events. Their use in modelling financial high-frequency data has picked up with many applications and variations.

In practice, however, time series entail changes in their dependence structure and the above mentioned models are more appropriate for stationary time series. It will be a crude approximation to adopt them in modeling the non-stationary processes without at least a minimal adaptation in a model's assumptions. There is also high risk in prediction and forecasting in doing so, as noted by Mercurio and Spokoiny (2004) for financial data. Dealing with high-frequency financial data is not less risky. A representative example is the U-shape in a stock's trading activity typically observed during a day. We cannot expect a market's behaviour to be like for like at every point of time. Information flows almost continuously and, therefore, the intraday dynamics vary significantly and cannot be ignored.

The simplest type of deviation from stationarity is, arguably, piecewise stationarity. This implies that the parameters of a stochastic process remain constant (hence, the process is stationary) for a certain period of time. In this paper, we focus on this type and we aim to identify the stationary segments by detecting the locations and number of change-points. If, of course, we knew the segments a-priori we would fit a stationary model to each of the segments and proceed with the prediction or forecasting tasks, but this knowledge is typically absent.

Detecting change-points has attracted significant attention. One approach to solve this problem is to formulate it through an optimization task i.e. minimising a multivariate cost function (or criterion) and adding a penalty when the number of change-points is unknown (see Yao 1988; Davis et al. 1995 among others). However, this approach typically comes with a high computational cost. To reduce the computational complexity, Killick et al. (2012) proposed PELT that has linear computational

time under the assumption that the number of change-points increases linearly with sample size.

An alternative approach to the change-point estimation problem is to minimize a series of univariate cost functions in a ‘greedy’ manner, i.e. detect a change-point and then progressively move to identify more. A popular representative of this category is the Binary Segmentation method (BS) and the reasons for its popularity are its low computational complexity and easiness of implementation: after identifying a single change-point (through the use of a certain statistic such as the CUSUM) the detection of further change-points continues to the left and to the right of the initial change-point until no further changes are detected.

The BS algorithm has been adopted to solve various types of problems with Inclan and Tiao (1994) to be, perhaps, the first to use it to detect breaks in the variance of a sequence of independent observations. For irregularly spaced time series, at least in the context of ACD or Hawkes processes, no literature exists that proposes a BS method (or an alternative change-point detection method) to detect change-points in a piecewise stationary ACD or Hawkes process. We note that Roueff et al. (2016) introduce the time-varying Hawkes process which is locally stationary, not necessarily piecewise stationary and, hence, they do not deal with change-point detection.

The BS algorithm, under specific model specifications, may be inefficient in detecting the change-points. Fryzlewicz (2014) illustrates that with a simple signal+noise model having only three change-points close to each other and in the middle of the series. At the beginning of the BS algorithm the CUSUM statistic does not clearly point to a change-point, hence, it does not move to search for more. This behaviour should not come as a surprise: BS is a greedy algorithm searching for a single change-point at each iteration and failing to do so results in an unnecessary early stop. The Wild Binary Segmentation algorithm (WBS; Fryzlewicz 2014) aims to solve the above mentioned limitation using a ‘certain random localization mechanism’. We discuss WBS later on, but we note here that this randomized algorithm was extended to univariate time series (Korkas and Fryzlewicz 2017) and high-dimensional panel data (Wang and Samworth 2018). Other variants of the random localization mechanism include the second version of WBS (WBS2) of Fryzlewicz (2020) and the Narrowest-Over-Threshold (NOT) method of Baranowski et al. (2019), while Anastasiou and Fryzlewicz (2019) propose Isolate-Detect (ID) which involves a fixed (rather than random) localization mechanism making it of an order of magnitude faster than the methods mentioned above. A faster implementation of WBS is SeedBS of Kovács et al. (2020) that relies on a deterministic construction of the localization mechanism.

In this work, we capitalize on BS’s popularity and propose a new randomised version of it which we term Ensemble Binary Segmentation (EBS). In particular, we draw a number M of random segments from a given univariate time series and apply the BS algorithm in each of these segments. The estimated change-points are then collected from over the M BS applications. Due to this simple mechanism the ways to combine the estimated change-points are numerous, for instance, the final set of change-points comprise of those that appear frequently or appear more frequently relative to other estimated change-points. An extra feature is that the estimated

change-points can be ranked based on their relative frequency which is useful when post-processing change-points.

The reader might question why WBS is not preferred for detecting change-points in a piecewise stationary ACD or Hawkes process and a new method is proposed instead. First, EBS exhibits better actual performance both in terms of computational speed and controlling the total number of change-points. The way that WBS works does not allow an efficient post-processing of the estimated change-points resulting in sometimes significant oversegmentation (in the case of a few change-points) or undersegmentation (in the case of frequent change-points). The source of the problem is the spurious detection of change-points which is a consequence of the distributional features of the ACD multiplicative form. Further, EBS, in practice, is fast because it does not search for a single change-point in every iteration, but rather it can locate all the change-points with even a single random draw with a small, albeit non-zero, probability. Second, contrary to WBS, the method we propose here can be extended to include other segmentation algorithms. In other words, when drawing a number of M of random segments we do not have to apply BS in each of these, but any other method can be potentially used. Therefore, studying EBS is a crucial starting point for other ensemble-type of change-point detection techniques.

The action of aggregating many method outputs to randomized versions of the data is not new. In fact, it has been studied extensively in the statistics and machine learning literature. Random forest is a popular representative that enjoys good predictive performance. Briefly, a random forest works as follows: grow multiple regression (decision) trees to randomized parts of the training dataset, and then average the trees. Our proposed method works in a similar manner and provides a fresh solution to the change-point problem.

The paper is structured as follows: in Sect. 2 we present the time-varying ACD model (tvACD), the core of our detection algorithm, and in Sect. 3 its transformation. The EBS algorithm is presented in Sect. 4, including its theoretical consistency in detecting the number and locations of change-points. The same section covers the post-processing of the estimated change-points and conducts simulation studies to obtain universal algorithm parameters that almost minimize the need for tuning by the user. Further, in Sect. 5 we conduct a simulation study to examine the actual performance of our proposed algorithm vis-à-vis the standard BS algorithm. In Sect. 6 we apply our method to financial high-frequency transaction data. Proofs of the results are in the supplementary material. Our methodology is implemented in the R package *eNchange* (Korkas 2020), available to download from CRAN.

2 Time-varying ACD model

Let τ_i be the transaction time where $0 = \tau_0 < \dots < \tau_{T-1}$ and $x_i = \tau_i - \tau_{i-1}$ be the duration between trades/events. We consider the following time-varying ACD model

$$x_i/\psi_i = \varepsilon_i \sim \mathcal{G}(\theta_1) \quad (1)$$

$$\psi_t = \omega(t) + \sum_{j=1}^p \alpha_j(t)x_{t-j} + \sum_{k=1}^q \beta_k(t)\psi_{t-k}. \quad (2)$$

where $\psi_t = \mathbb{E}[x_t | x_{t-1}, \dots, x_{t-p}, \psi_{t-1}, \dots, \psi_{t-q}; \Theta(t)]$ is the conditional mean duration of the t -th event with parameter vector $\Theta(t)$ and $\mathcal{G}(\cdot)$ is a general distribution over $(0, +\infty)$ with mean equal to 1 and parameter vector θ_1 . The vector of parameters involved in the conditional mean ψ_t is

$$\Theta(t) = (\omega(t), \alpha_j(t), \beta_j(t)) \in \mathbb{R}^{1+p+q}.$$

Then, $\Theta(t)$ is piecewise constant in t with N change-points η_i ($\eta_0 = 0, \eta_{N+1} = T$) i.e. at any η_i we have that $\Theta(\eta_i) \neq \Theta(\eta_i + 1)$ for all $i = 1, \dots, N$. The number of change-points and their locations are allowed to increase with T , hence, formally, we have that $N = N(T)$ and $\eta_i = \eta_i(T)$ for $\forall i = 1, \dots, N$. To save notation and be consistent with the existing literature on change-point detection, henceforth, we use the short-hand notation N and η_i .

The number of change-points N and their locations are assumed to be unknown and we aim to estimate them. The parameter values in $\Theta(t)$ are also unknown, but these can be estimated after the change-points have been detected and stationary segments are identified.

We further assume the following conditions on (1) and (2)

(A1) For some $\epsilon_1 > 0, \xi_1 < \infty$ and all T , we have

$$\inf_{1 \leq t \leq T} \omega(t) > \epsilon_1 \text{ and } \sup_{1 \leq t \leq T} \omega(t) \leq \xi_1 < \infty.$$

(A2) For some $\epsilon_2 \in (0, 1)$ and all T , we have

$$\sup_{1 \leq t \leq T} \left\{ \sum_{j=1}^p \alpha_j(t) + \sum_{k=1}^q \beta_k(t) \right\} \leq 1 - \epsilon_2.$$

Assumptions (A1)–(A2) guarantee that between any two consecutive change-points, x_t admits a well-defined solution a.s. and is weakly stationary (see e.g., Theorem 4.35 of Douc et al. 2014).

For a non-stationary stochastic process x_t , its strong-mixing rate is defined as a sequence of coefficients $\alpha(\kappa)$ such that

$$\alpha(\kappa) = \sup_{t \in \mathbb{Z}} \sup_{\substack{G \in \sigma(X_{t+\kappa}, X_{t+\kappa+1}, \dots), \\ H \in \sigma(X_t, X_{t-1}, \dots)}} |\mathbb{P}(G \cap H) - \mathbb{P}(G)\mathbb{P}(H)|. \quad (3)$$

In Fryzlewicz and Subba Rao (2011) the mixing rate of univariate, time-varying ARCH processes was investigated, and in Cho and Korkas (2018) the mixing rate of any pair of time-varying GARCH processes was established. A tvACD process is also strong mixing at a geometric rate, under the following Lipschitz-type condition on the density f_{ϵ_t} of ϵ_t .

(A3) The distribution of ε_t satisfies the following: for any $a > 0$, there exists fixed $K > 0$ independent of a such that

$$\int |f_{\varepsilon_t}(u) - f_{\varepsilon_t}(u[1+a])| du \leq Ka.$$

Fryzlewicz and Subba Rao (2011) show that (A3) is satisfied when certain conditions about a density function f hold i.e. the first derivative is bounded; after some finite point m the derivative f' declines monotonically to zero; and $\int_0^\infty |zf'| dz < \infty$. It is straightforward to show that f_{ε_t} satisfies these conditions, hence, (A3) holds when $\varepsilon_t \sim \exp(\theta_1)$. Many well-known distributions satisfy this condition (see Fryzlewicz and Subba Rao 2011), but extension of our method to these will require a transformation of the duration process that accounts for the various distribution parameters. This is a technically challenging task, and it is left for future research. In this work, we assume that distribution $\mathcal{G}(\theta_1)$ in (1) is $\exp(1)$.

A different approach to modeling the arrival times is through the Hawkes process: a self-exciting process where the arrival of an event (or events) causes the conditional intensity function $\lambda(\tau)$ to increase. In this work, we introduce the time-varying Hawkes process (which we term *tvHawkes*). In particular, we consider a counting process on $\mathbb{R}^+ \{\mathcal{N}_\tau\}_{\tau \geq 0}$ with associated history $\{\mathcal{F}_\tau\}_{\tau \geq 0}$ whose conditional intensity is defined as

$$\lambda(\tau) = \lambda_0(\tau) + \sum_{\tau_i < \tau} \sum_{j=1}^p \alpha_j(\tau) e^{-\beta_j(\tau)(\tau - \tau_i)}, \quad \text{for } \tau = 1, \dots, T^h. \quad (4)$$

where $\lambda_0 > 0$ is the initial intensity at time $\tau = 0$ and $\alpha_j(\tau)$, $\beta_j(\tau)$ are positive parameters that are piecewise constant in τ with N change-points η_i . We refer the reader to the supplementary material for more details.

To prove consistency of EBS in detecting the number of change-points and their locations when the underlying model is a tvHawkes process we would need assumptions around non-negativity of $(\lambda_0(\tau), \alpha_j(\tau), \beta_j(\tau))$ and stationarity in between any two consecutive change-points. However, to the best of our knowledge, no mixing rates have been established for a time-varying Hawkes process and, therefore, it is not trivial to estimate an upper bound for the cumulative error term. In this work section, we mainly focus on tvACD while we use tvHawkes as a benchmark (and an alternative) to tvACD.

3 Transformation of the duration process

In the first stage of our proposed methodology we form the process U_t using a function $g_0: \mathbb{R}^{1+p+q} \rightarrow \mathbb{R}$ that takes $\mathbf{x}_t^{1-p} = (x_t, x_{t-1}, \dots, x_{t-p})^T$ and $\boldsymbol{\psi}_{t-1}^{1-q} = (\psi_{t-1}, \dots, \psi_{t-q})^T$ as inputs. This function g_0 is required to be bounded and Lipschitz continuous.

(A4) The function $g_0 : \mathbb{R}^{1+p+q} \rightarrow \mathbb{R}$ satisfies $|g_0| \leq \bar{g} < \infty$ and is Lipschitz continuous, i.e., $|g_0(z_0, \dots, z_{p+q}) - g_0(z'_0, \dots, z'_{p+q})| \leq C_g \sum_{k=0}^{p+q} |z_k - z'_k|$.

Empirical residuals are widely adopted for detecting changes in the parameters of a stochastic model and in this work it will form the basic statistic for our detection algorithm in the context of point processes.

For the tvACD model, following Fryzlewicz and Subba Rao (2014) we select g_0 such that

$$U_t = g_0(\mathbf{x}_t^{t-p}, \boldsymbol{\psi}_{t-1}^{t-q}) = \frac{x_t}{\tilde{\psi}_t}, \quad \text{where} \quad (5)$$

$$\tilde{\psi}_t = C_0 + \sum_{j=1}^p C_j x_{t-j} + \sum_{k=1}^q C_{p+k} \psi_{t-k} + \varepsilon x_t$$

where the last term εx_t is added to ensure the boundness of U_t . This transformation decorrelates the original tvACD process and lightens its tails. Therefore, it serves as our main change-point detection statistic by observing that when any parameter $\Theta(t)$ undergoes a change at some point t , so does $\mathbb{E}(U_t)$. In this work we favour $\log(U_t + \epsilon)$ where ϵ as in (5) to reduce the rightward skew observed in the distribution of U_t . The log-transformation does not impact the mixing properties of the process and, as a result, our consistency result for U_t still holds for $\log(U_t + \epsilon)$ (see, e.g., Theorem 14.1 in Davidson 1994).

It is important to note that x_t or any ‘diagonal’ transformation of x_t such as $\log(x_t^2)$ should not be used instead of U_t because both x_t and the logarithmic process $\log(x_t^2)$ are highly autorrelated. This will distort the change-point detection most likely by producing a false picture of the locations of the change-points. For a discussion on transformations we refer to the original work of Fryzlewicz and Subba Rao (2014) and for the tvHawkes transformation to the supplementary material.

We prepare the ground for a consistency result for our proposed method. Let $\{bx_t^i\}_{t=1}^T$ denote a stationary ACD process with parameters $\Theta(\eta_i + 1)$, and the innovations coinciding with ε_t over the associated segment $[\eta_i + 1, \eta_{i+1}]$. For each $i = 1, \dots, N$ we also define $\tilde{U}_t^i = g_0(\mathbf{x}_t^{i,t-p}, \boldsymbol{\psi}_{t-1}^{i,t-q})$ and denoting the index of the change-point strictly to the left and nearest to t by $v(t) = \max\{0 \leq i \leq N : \eta_i < t\}$ with which $\{x_t^i\}_{t=1}^T$ and $\tilde{U}_t^i$ are defined.

Proposition 1 Suppose that (A1)–(A4) hold, and let $y_t = U_t$. Then, we have the following decomposition

$$y_t = f_t + z_t, \quad 1 \leq t \leq T. \quad (6)$$

- (i) f_t are piecewise constant as $f_t = \tilde{g}_t$ where $\tilde{g}_t = \mathbb{E}\{(\tilde{U}_t^{v(t)})\}$. All change-points in f_t belong to $\mathcal{B} = \{\eta_1, \dots, \eta_N\}$.
- (ii) z_t satisfies

$$\max_{1 \leq s < e \leq T} \frac{1}{\sqrt{e-s+1}} \left| \sum_{t=s}^e z_t \right| = O_p(\sqrt{\log T}).$$

Proof of Proposition 1 can be found in the supplementary material. Unlike $\mathbb{E}(U_t)$, we have $\tilde{g}_t$ which is *exactly* constant between any two adjacent change-points without any boundary effects. By its construction, $z_t = U_t - \mathbb{E}(\tilde{U}_t^{v(t)})$ does not satisfy $\mathbb{E}(z_t) = 0$, but thanks to the mixing properties of U_t as in (3), its scaled partial sums can be appropriately bounded. In the next section, we introduce the multiple change-point detection algorithm.

4 The segmentation algorithm

4.1 The binary segmentation algorithm

We first present the binary segmentation algorithm within the framework of the ACD model. In the next section, we explain how this algorithm can be enhanced by ensembling detected change-point from multiple applications of the BS algorithm on smaller (and random) segments.

We consider the CUSUM-type statistic which has been widely adopted for change-point detection in both univariate and multivariate data

$$\mathbb{Y}_{s,e}^b = \sqrt{\frac{(b-s+1)(e-b)}{e-s+1}} \left(\frac{1}{b-s+1} \sum_{t=s}^b y_t - \frac{1}{e-b} \sum_{t=b+1}^e y_t \right) \quad (7)$$

for $s \leq b < e$. A large value of $|\mathbb{Y}_{s,e}^b|$ typically indicates the presence of a change-point in the level of y_t in the vicinity of $t = b$. In particular, if $\mathbb{Y}_{s,e} > \pi_T$ where

$$\mathbb{Y}_{s,e} = \max_{s \leq b < e} |\mathbb{Y}_{s,e}^b|, \quad (8)$$

and π_T is a threshold (the choice of which is discussed in Sect. 4.6), then the location of the change-point is estimated as

$$\hat{\eta} = \arg \max_{s \leq b < e} \mathbb{Y}_{s,e}^b.$$

The BS algorithm starts by initialising with $s = 1$, $e = T$ and $\hat{\mathcal{B}} = \emptyset$ and proceeds by recursively applying the CUSUM statistic (7) on $[s, \hat{\eta}]$ and $[\hat{\eta} + 1, e]$. The BS algorithm stops in each current interval when no further change-points are detected, the obtained CUSUM values fall below threshold π_T .

4.2 The ensemble binary segmentation algorithm

Our recommended enhancement of the BS algorithm, which we argue to be a significant improvement, is based on the fact that the BS method can possibly fit the wrong

model when multiple change-points are present as it searches the whole series. The application of the CUSUM statistic can result in spurious change-point detection when e.g. the true change-points occur close to each other. Especially, Olshen et al. (2004) notice that BS method can fail to detect a small change in the middle of a large segment which is illustrated in Fryzlewicz (2014).

To solve this, Fryzlewicz (2014) proposes a randomised version of the binary segmentation method (termed Wild Binary Segmentation—WBS) where the search for change-points proceeds by calculating the CUSUM statistic in smaller segments whose length is random. By doing so, WBS aims to draw favourable intervals with probability tending to one with the sample size containing at most a single change-point. To improve upon WBS mainly in reducing complexity, Fryzlewicz (2020) suggests a new recursive algorithm, termed wild binary segmentation 2 (WBS2), for producing a ‘complete’ solution path i.e. a sequence of estimated nested models containing $0, \dots, T - 1$ change-points. To achieve this he computes all possible CUSUM statistics in any given interval when it is not too long; otherwise it computes the CUSUM statistics in random segments as in WBS. The final selection of the change-points is done through the ‘Steepest Drop to Low Levels’ procedure that is not penalty-based, and only uses thresholding as a certain discrete secondary check (WBS2.SDLL). Other solutions to the BS issue described above are that of Baranowski et al. (2019) with NOT and Anastasiou and Fryzlewicz (2019) with ID.

We provide reasons for choosing EBS over WBS or WBS2. First, both WBS and WBS2 are tailored around BS, and it is not always straightforward to apply to other segmentation methods. On the other hand, EBS can be easily extended as it collects estimated change-points from applying BS to many random sub-samples. In addition, it is typical to oversegment a series under the multiplicative setting which in turn requires further processing of the estimated change-points. How well this post-processing will perform depends on the detection ability of WBS itself both in localisation accuracy and the total number of change-points identified. EBS, as we explain shortly after, is naturally equipped with a ‘second chance’ mechanism of selecting the change-points and is less likely to oversegment, which gives it an advantage before entering the post-processing stage. Because oversegmentation does not arise in the signal+iid Gaussian noise setting of Fryzlewicz (2014), as it is entirely due to the distributional features of the multiplicative setting where the input sequence is a rightward skewed and autocorrelated random sequence, the complete solution path produced by WBS2.SDLL will not be much of use in post-processing either. The reason is that in order to rank the change-points estimated from WBS2.SDLL we would need to rank the CUSUM statistics, which in turn can take large values when the starting and ending points are close to a change-point. A solution could be to adjust for a segment’s (random) length or average the CUSUM statistics for each change-point (again after adjusting for the random lengths of the segments), however, we favour EBS for being independent of the location statistic when ranking the change-points.

We propose a randomised algorithm, but instead of drawing random intervals, calculate the CUSUM statistic in each of these and proceed in a “to-the-left-and-to-the-right” manner, we run M multiple BS on random segments of the underlying univariate series. We then collect all the obtained change-points and calculate

the frequency of the occurrences by simply counting the number of times a certain change-point appears over M draws of the BS algorithm. This results in a better performance compared with BS as it is more likely to draw favourable intervals which can contain a single change-point (as in the WBS methodology) or more than one (as in the BS methodology). By doing so we aim to balance the benefits of both worlds. The action of ensembling (hence, borrowing from the machine learning literature we term our methodology Ensemble BS or EBS) allow us, through a slightly randomised mechanism, to classify a change-point as important when it appears more often than other change-points.

In the last stage, we have the options to either (1) inspect the histogram of the estimated change-points; or (2) to keep only the change-points that appear more than $\lceil \pi_{thr} \cdot M \rceil$ times where $\pi_{thr} \in [0, 1]$; or (3) to apply (2) and then post-process the detected change-points by removing change-points that are ranked lower and/or are ‘close’ to high ranked change-points.

As an illustration, we simulated a tvACD model of the following form

$$x_t = (\omega(t) + 0.1x_{t-1} + 0.7\psi_{t-1}) \cdot \varepsilon_t \sim \exp(1) \quad (9)$$

with sample size $T = 3000$ and $\omega(t) \in \{1/16, 1/4, 1/16, 1/4, 1/16\}$ at $.475T, .485T, .495T$, and $.505T$.

We chose this setup because change-points in the middle of a long time series is challenging for BS to detect. The CUSUM statistic (in absolute value) will fail to exceed the threshold and the BS algorithm will stop. On the contrary, when taking random intervals (s_m, e_m) and then calculating the CUSUM statistic it was more likely for it to obtain a high CUSUM statistic in absolute values (hence, more likely to exceed the threshold), see the right top plot in Fig. 1. What is important to observe is that only s_m has to be close to the first change-point in order for the CUSUM statistic to surpass the threshold, while end point e_m can be much further to the right (observe the blue area from 1200 onwards in the x-axis). This observation holds symmetrically i.e. if e_m is close to the last (fourth) change-point then s_m can freely take any value in the interval $[1, [0.475T]]$ (observe the blue area below 1600 in the y-axis). In those both favourable scenarios (blue areas) the BS algorithm will proceed to identify the rest change-points consistently.

It is interesting to note that the number of draws M did not alter the change-point performance. From the bottom two plots in Figure 1 one can see that the shapes of the empirical distributions look identical even though in the first case $M = 1000$, while in the second M was 10-fold. The reason for that lies in the proof of the consistency result (see the supplementary material), but it can be easily seen by inspecting the model (9): a favourable draw can occur when *only* the starting point s is in a local neighbourhood to the first change-point from the left, and applying a single BS routine will result in to identify all adjacent change-points. Contrary to WBS, whereby it would seek to localise again after the identification of a change-point, EBS would likely identify all the change-points from the first (favourable) draw.

We now describe the EBS algorithm in more detail. First, denote by $\mathcal{D}_m = \{\hat{\eta}_1^{(m)}, \hat{\eta}_2^{(m)}, \dots\}$ for $m = 1, \dots, M$ the set of all the change-points detected by the BS algorithm on a random interval $[s_m, e_m] \subseteq [1, T]$. Let $\mathcal{E} = \bigcup_{m=1}^M \mathcal{D}_m$, the

Fig. 1 A simulated tvACD process from the model (9) (top left). A heatmap of the maximum of the absolute CUSUM statistics calculated on the transformed series U_t with varying start (s_m on x-axis) and end (e_m on y-axis) points (top right). The empirical histogram of the estimated change-points from $M = 1000$ draws (bottom left). The empirical histogram of the estimated change-points from $M = 10,000$ draws (bottom right). The four vertical bars in red are the real locations of the change-points

ensemble collection of all the detected change-points from the M applications of BS, and $|\mathcal{E}|$ its cardinality. Evaluate each change-point $\hat{\eta}_i$ by counting the number of times (frequency) it appears in the ensemble $\mathcal{E}$, i.e.

$$V_{\hat{\eta}_i}^{(M)} = \sum_{m=1}^M \mathbb{I}(\hat{\eta}_i \in \mathcal{D}_m) \text{ for } i = 1, \dots, \hat{N} \quad (10)$$

where $\mathbb{I}(\cdot)$ is the indicator function. In the machine learning literature, formula (10) is referred as majority voting. We can also obtain the relative frequency of a change-point η_i using $\hat{f}_{\hat{\eta}_i} = V_{\hat{\eta}_i} / |\mathcal{E}|$ to create a relative importance plot (histogram). To make a decision about how many change-points are finally selected (and, hence, omit falsely detected change-points), one way is to select from the ensembles that $\hat{\eta}_i$ such that

$$V_{\hat{\eta}_i} \geq \pi_{thr} \cdot M, \quad \pi_{thr} \in [0, 1]. \quad (11)$$

Ensembling the change-points from a random interval sampling scheme is obviously not the only way to go. From a theoretical (and practical) point of view, EBS can be based on a deterministic grid of intervals as, for example, in Kovács et al. (2020) whose segmentation method leads to a near-linear time approach. A valid argument here is that a random scheme is likely not appropriate because of the overlapping random intervals. However, it is only with a random interval sampling scheme that we can get a meaningful ‘aggregation’ result, whereby the frequency of a detected change-point across M runs can be seen as a measure of its importance.

We emphasize that other types of thresholding can be utilized, for instance, a change-point is selected when it is ranked above the average, i.e.

$$V_{\hat{\eta}_i} > \frac{1}{\hat{N}} \sum_{k'=1}^{\hat{N}} V_{\hat{\eta}_{k'}} \quad \forall k' = 1, \dots, \hat{N}.$$

The relative importance plot also provides a type of ‘scree plot’, whereby an elbow (kink) indicates what the right number of selected change-points is. This type of scree plot is common in, for example, principal component analysis and the selection of the number of principal components.

In this work, we opt for the majority voting type of selection and we point out that the selection based on ranking lays ground for further research. We explain why by looking at the problem of change-point detection through the lens of variable selection (see, e.g., Harchaoui and Lévy-Leduc 2010). We identified two views in randomised versions of variable selection. Meinshausen and Bühlmann (2010) argue that a variable is declared predominant if a slightly randomized routine selects it ‘more frequently’ than other variables (i.e. majority votes), and a threshold parameter exists that controls exactly what we mean by ‘more frequently’. However, Xin and Zhu (2012) argue that variable ranking is more practical than variable selection as such, and the ability to rank variables based on how often they are selected is a major advantage of the ensemble approach making it more attractive than other selection algorithms (see also Zhu 2015 for an interesting review of majority voting in variable selection). Finally, to avoid confusion, we stress that a change-point that ranks high within $\mathcal{E}$ should be preferred over other change-points, but only (and optionally) for post-processing the change-points.

Algorithm 1: EnBinSeg (Ensemble Binary Segmentation algorithm)

Input: $\{y_t\}, \pi_T, \pi_{thr}, \hat{\mathcal{E}}, \hat{\mathcal{B}}, M$

Step 1: for $1 \dots M$ do

$[s_m, e_m] \leftarrow U(1, T)$
 $\hat{\mathcal{B}}^{(m)} \leftarrow \text{BinSeg}(\{y_{t \in s_m:e_m}\}, \pi_T, s_m, e_m, \hat{\mathcal{E}})$
 $\hat{\mathcal{E}} \leftarrow \hat{\mathcal{E}} \cup \hat{\mathcal{B}}^{(m)}$

end

Step 2: for $i = 1 \dots \hat{N}$ do

if $\hat{\eta}_i \in \hat{\mathcal{B}}^{(m)}$ then
 $v_{m, \hat{\eta}_i} \leftarrow 1$
 else
 $v_{m, \hat{\eta}_i} \leftarrow 0$
 end

end

Step 3: for $i = 1 \dots \hat{N}$ do

$V_{\hat{\eta}_i} = \sum_{m=1}^M v_{m, \hat{\eta}_i}$
 if $V_{\hat{\eta}_i} > \pi_{thr} \cdot M$ then
 $v_{m, \hat{\eta}_i} \leftarrow 1$
 $\hat{\mathcal{B}} \leftarrow \hat{\eta}_i$

end

Step 4 (Optional for post-processing): Rank $\hat{\eta}_i$ by $V_{\hat{\eta}_i}$ for

$i = 1, \dots, \hat{N}$

Output: $\hat{\mathcal{B}}$

4.3 Theoretical properties of EBS

In this section we present the consistency theorem for the EBS algorithm for the total number N and locations of the change-points $0 < \eta_1 < \dots < \eta_N < T - 1$ with $\eta_0 = 0$ and $\eta_{N+1} = T$. To achieve this, we impose the following conditions in addition to (A1) to (A4) mainly to control the detectability of each η_b .

- (B1) The number of change-points N in (1)–(2) is unknown and allowed to increase with T and only the minimum distance between the change-points can restrict the maximum number of N .
- (B2) There exists a fixed constant $f^* > 0$ such that $\max_{1 \leq t \leq T} |f_t| \leq f^*$.
- (B3) The distance between any two adjacent change-points satisfies $\min_{r=1, \dots, N+1} |\eta_r - \eta_{r-1}| \geq \delta_T$, where $\delta_T \geq C \log T$ for a large enough C .
- (B4) The magnitudes of the change-points satisfy $\inf_{1 \leq i \leq N} |f_{\eta_i+1} - f_{\eta_i}| \geq f_\star$ where $f_\star > 0$.

Theorem 1 Suppose that Assumptions (A1)–(A4) and (B1)–(B4) hold. With the number of change-points as N and the locations of those change-points as $\eta_1, \dots, \eta_N$, let $\hat{N}$ and $\hat{\eta}_1, \dots, \hat{\eta}_N$ be the number and locations of the change-points (in ascending order) estimated by the Ensemble Binary Segmentation algorithm. There exist constants C_1 and C_2 such that if $C_1 \sqrt{\log T} \leq \pi_T \leq C_2 \sqrt{\delta_T}$, then $P(\mathcal{Z}_T) \geq 1 - \pi_z - (1 - \delta_T T^{-1}/9)^M$, where

$$\mathcal{Z}_T = \{\hat{N} = N; \max_{i=1, \dots, N} |\hat{\eta}_i - \eta_i| \leq C \log T\}$$

for certain C and $\pi_z = \bar{C}_1 T^2 \exp(-\bar{C}_2 (c' \sqrt{\log T} - \bar{C}_3 (p+q) N \log(T)^{-1/2})^2)$ for certain $c', \bar{C}_1, \bar{C}_2$ and $\bar{C}_3$. The guaranteed speed of convergence of $P(\mathcal{Z}_T)$ to 1 is no faster than $(1 - \delta_T T^{-1}/9)^M$ where M is the number of random draws.

The rate of convergence for the estimated change-points obtained for the BS method by Fryzlewicz and Subba Rao (2014) is $\mathcal{O}(\sqrt{T} \log^{\bar{\alpha}} T)$ where $\bar{\alpha} > 0$ when δ_T is of order T . In the EBS setting, the rate is square logarithmic when δ_T is of order $\log T$, which represents an improvement.

We now elaborate on the minimum number M of random draws required to ensure that the bound on the speed of convergence of $P(\mathcal{Z}_T)$ to 1 in Theorem 1 is suitably small. Suppose that we wish to ensure that

$$(1 - \delta_T T^{-1}/9)^M \leq T^{-1}.$$

This is equivalent to

$$M \geq \frac{9T}{\delta_T} \log(T)$$

by noting that $\log(1 - y) \approx -y$ around $y = 0$.

Let us consider the “easiest” case, i.e. $\delta_T \sim T$. This results in a logarithmic number of draws, which leads to particularly low computational complexity, and it also has the same complexity with the WBS case. When $\delta_T \sim \sqrt{\log(T)}$, then the required M increases almost linearly, but it has less computational complexity than WBS, which also explains why EBS is generally faster than WBS.

Finally, we discuss π_{thr} which appears in (11), and it acts as a decision rule when aggregating estimated change-points across M applications of the BS algorithm. In practice, EBS tended to return spurious change-points due to the distributional features of the multiplicative setting which require us to control the partial sums of z_t in (ii) of Proposition 2 achieved through an appropriately chosen π_{thr} . Theoretically, it is not trivial to control the falsely detected change-points, as the distribution of the partial sums of z_t is not easy to obtain, and we have to rely on simulations to identify an appropriate choice of π_{thr} . Our practical recommendations for the choice of π_{thr} along with the choice of M are discussed in Sect. 4.6.

4.4 Post-processing

In practice, the real number of change-points is not known to us and to reduce the risk of over-segmentation we propose a post-processing method. The need for post-processing the estimated change-points from a detection routine is common within the context of multiplicative models. We refer the reader to Inclan and Tiao (1994), Cho and Fryzlewicz (2012) and Korkas and Fryzlewicz (2017). In these works, the post-processing method compares every detected change-point from the main detection method against the adjacent ones (re)using the CUSUM statistic. Even though there are variations in implementing a post-processing between them, their common drawback is the lack of information about the importance of the detected change-points.

Our proposal in filtering estimated change-points is simple: starting with the highest ranked $\hat{\eta}_i^{(1)}$ based on its $f_{\hat{\eta}_i}$, we remove from set $\hat{\mathcal{B}} \hat{\eta}_i^{(2)}$ if it is within a distance Δ_T and examine whether $\hat{\eta}_i^{(3)}$ is within this distance. If $\hat{\eta}_i^{(2)}$ is not within distance, we keep this change-point and repeat the process until all change-points are separated by at least Δ_T . Because Δ_T cannot be smaller than the minimum permitted distance δ_T —as in Assumption (B3)—between any two adjacent change-points, hence, $\Delta_T = \mathcal{O}(\delta_T)$. In practice, we take $\Delta_T = \lceil 0.005 \cdot T \rceil$.

4.5 Choice of parameters for transformation

The right choice of the transformation function g_0 will influence the empirical performance of our methodology and its power in particular. This choice boils down to determine the coefficients C_j , $j = 0, \dots, p + q$.

The BASTA-res algorithm of Fryzlewicz and Subba Rao (2014) performs change-point detection in the univariate ARCH process, as already seen similar to the ACD process, by analysing the transformation of the input time series obtained similarly to U_t^2 . They recommend the use of ‘dampened’ versions of the GARCH

parameter estimates which we also adopt here for the ACD process. This leads to the choice of $C_0 = \hat{\omega}$, $C_j = \hat{\alpha}_j/F$, $1 \leq j \leq p$ and $C_{p+k} = \hat{\beta}_k/F$, $1 \leq k \leq q$, with $F \geq 1$.

Empirically, the motivation behind the introduction of F is as follows. For x_t with time-varying parameters, we often observe that α_j and β_k are over-estimated in the sense that $\sum_{j=1}^p \hat{\alpha}_j + \sum_{k=1}^q \hat{\beta}_k$ is close to one, especially when dealing with real data. There has been evidence in the literature that change-points may cause persistence estimation in volatility models (e.g. Francq et al. 2001). Mikosch and Stărică (2004), among others, show that the estimated persistence close to unity in GARCH models is likely spurious and confounded by neglected change-points. We observed the same phenomenon when trying to fit an ACD process to simulated and real data. Using the raw estimates in place of C_j 's in (5), therefore, is not the best approach. Cho and Korkas (2018) propose to choose the dampening factor F as

$$F = \max \left[1, \frac{\min(0.99, \sum_{j=1}^p \hat{\alpha}_j + \sum_{k=1}^q \hat{\beta}_k)}{\max \{0.01, 1 - (\sum_{j=1}^p \hat{\alpha}_j + \sum_{k=1}^q \hat{\beta}_k)\}} \right]. \quad (12)$$

By construction, F is bounded as $F \in [1, 99]$ and approximately brings $\hat{\omega}$ and $\sum_{j=1}^p \hat{\alpha}_j + \sum_{k=1}^q \hat{\beta}_k$ to the same scale. The selection of the order p, q can be done through the means of an information criterion; however, taking into consideration the possibility that the estimated persistence will be close to unity, we conducted a simulation experiment to establish the right choice of order. For brevity, the details of the experiment are provided in the supplementary material.

The results indicate that the transformation involving the empirical estimates from an ACD(0,1) model vis-a-vis ACD(1,1) had a better performance. The dampening factor F did not have as a strong impact and generally the performance remained unchanged for increasing F . In a separate experiment not shown here the dynamic selection of F worked better and is, hence, recommended.

4.6 Choice of threshold and parameters

In this section we present choices of the parameters involved in the main Theorem 1. In particular, we have that the threshold π_T includes the constant C_1 . To approximate the distribution of the CUSUM statistic in the absence of change-points one approach is to use a parametric resampling procedure that was described in Cho and Korkas (2018) and which provides good results. A similar approach has widely been adopted in the change-point literature including Kokoszka and Teyssière (2002) who test for the presence of a change-point in the parameters of univariate GARCH models. However, this resampling procedure adds computational time and it does not fully take advantage of the dampening transformation aiming to bring a series closer to an iid exponential distribution. For that reason we conduct experiments to establish the *universal* value of the threshold parameter under the null hypothesis of no change-points, so that the threshold π_T overall depends only on the sample size T .

In particular, we generate stationary ACD processes of size T , varying from 500 to 100,000 with a step 50. The exact choice of the ACD model parameters (i.e.

$(\omega, \alpha, \beta))$ did not alter the results and, therefore, not reported here. Then we find b that maximises (7). The ratio

$$C_1 = \mathbb{Y}_{s,e}^b (\sqrt{\log T})^{-1}$$

gives us an insight into the magnitude of parameter C_1 . We repeated this experiment 100 times for different values of T and we selected C_1 as the $100(1 - \alpha)\%$ -percentile for each instance of T . Our results indicated that C_1 tends to decrease as we increase the sample size and remains unchanged after a certain point (see Figure 2 in the supplementary material).

To propose a general rule that will apply in most cases we fitted the regression

$$C_1(T) = c_0 + c_1 T + c_2 \frac{1}{T} + c_3 T^2 + \varepsilon.$$

Having estimated the values for $\hat{c}_0, \hat{c}_1, \hat{c}_2, \hat{c}_3$ we were able to use fitted values for any sample size T . For samples larger than $T = 100,000$, we used the same C_1 values as for $T = 100,000$.

We turn to the choice of M —the number of ensembles to run - and π_{thr} —the relative frequency threshold. In Sect. 4.3 we discussed that the minimum number of M is in the range of a few thousands when the distance of any two adjacent change-points is of $O((\log T)^{1/2})$. However, in practice and thanks to the randomised ensemble mechanism of EBS, M can be much smaller. As for π_{thr} , the theoretical choice of 0 is adequate, albeit for the distributional features of the multiplicative setting we consider here a non-zero positive value can work better in practice. Similar to the procedure of the choice of C_l above, we simulate stationary ACD processes and we apply the EBS algorithm for every triplet (M, π_{thr}, T) keeping everything else the same; see the supplementary material for more details. By inspecting Figure 4 in the supplementary material, we can see that the choice of $\pi_{thr} = 0.05$ results in a significant improvement in the detection accuracy compared with lower values. For larger values, we do not see any further improvement, even when the sample size increased. In addition we conclude that, conditioning on π_{thr} , a high number of draws M did not result a considerable improvement in accuracy. On the contrary, a low number in the range of hundreds gave similar results to $M = 5000$ even for larger samples. For the simulation study and the real application we set $\pi_{thr} = 0.05$ and $M = 500$ which are also the default values in the R package.

5 Simulation study

We simulated stationary time series with innovations $\varepsilon_t \sim \exp(1)$ for ACD and sample size $T = 2000$ for different specifications. For the Hawkes process we chose the time horizon $h = 500$, whereby the sample size will depend on the choice of the parameters. Roughly speaking, for $\lambda_0 = 1$, $\alpha_1 = 0.1$ and $\beta_1 = 0.7$ the sample size $T \approx 5000$. We report the false positive rate for each method i.e. the percentage of incorrectly rejecting the null hypothesis of no change-points.

From Table 1, we see that both BS and EBS performed well meaning that the risk of segmenting a stationary process is generally limited. WBS displayed oversegmentation that became more significant such as in cases of processes with high persistence, that is when $\alpha_1 + \beta_1$ in the ACD process or the ratio α_1/β_1 in the Hawkes process is either approaching 1. For example, in the rather challenging case of S3, EBS performed similar to BS, while WBS returned a high false positive ratio (in 68% of the cases incorrectly rejected the null hypothesis). It is worth mentioning (not shown here) that the risk of oversegmentation increased for WBS when a higher number of random draws was chosen, while the false positive rate for EBS remained unchanged for a larger M (1500 instead of the default 500).

We now examine the detection performance of our method for a set of non-stationary models, both ACD and Hawkes processes. We consider various test models and examine BS and EBS performance over 100 simulations for each of the test model. To compare the performance between the two detection methods we calculate the following error measures: $E(\hat{N} - N)$, $E(|\hat{N} - N|)$, $E[(\hat{N} - N)^2]$.

Since EBS has improved rates of convergence compared with BS and it can also assist in the post-processing stage, its accuracy should be judged in parallel with the total number of change-points identified. We use a test from Korkas and Fryzlewicz (2017) that tries to accomplish this. Assuming that the maximum distance from a real change-point η is denoted by $d_{\max}$, an estimated change-point $\hat{\eta}$ is correctly identified if $|\eta - \hat{\eta}| \leq d_{\max}$ (here within 1% of the sample size). In the case where two (or more) estimated change-points are within this distance $d_{\max}$ then only one change-point, the closest to the real change-point, is classified as correct and the rest are deemed to be false, except if any of these are close to another change-point. An estimator performs well when the hit ratio

$$HR = \frac{\text{\#correct change-points identified}}{\max(N, \hat{N})}$$

Table 1 Stationary processes results

Model	BS	WBS	EBS
S1: iid standard poisson with parameters $\lambda_0 = 2$ ($\alpha_1 = \beta_1 = 0 \forall j, k$) and $T = 2000$	2% (0.018)	18% (0.119)	4% (0.061)
S2: Hawkes process with parameters $\lambda_0 = 0.5$, $\alpha_1 = 0.1$ and $\beta_1 = 0.7$ ($T_{\text{avg}} = 250$)	0% (0.013)	0% (0.033)	2% (0.045)
S3: Hawkes process with parameters $\lambda_0 = 5$, $\alpha_1 = 0.4$ and $\beta_1 = 0.7$ ($T_{\text{avg}} = 5825$)	22% (0.337)	68% (1.158)	26% (0.415)
S4: ACD process with parameters $\omega = 1$, $\alpha_1 = 0.1$ and $\beta_1 = 0.7$ and $T = 2000$	4% (0.017)	25% (0.123)	8% (0.066)
S5: ACD process with parameters $\omega = 3$, $\alpha_1 = 0.15$ and $\beta_1 = 0.5$ and $T = 2000$	9% (0.180)	29% (0.161)	9% (0.063)

For each of the Hawkes processes the average sample size T_{avg} is given in brackets. Figures show the false positive rate for each method. Average execution time in seconds is given in brackets

is close to 1. According to the authors, the term $\max(N, \hat{N})$ penalises cases where, for example, the estimator correctly identifies a certain number of change-points all within the distance $d_{\max}$, but $\hat{N} < N$. It also penalises the estimator when overestimates the total number of change-points and all $\hat{N}$ detected change-points are within distance $d_{\max}$ of the true ones.

We consider nine model specifications M1–M8.2 ranging from a model with a single change-point to models with 19 change-points (multi-teeth types similar to Fryzlewicz (2014) for the additive signal + noise setups) and to models with random settings for the number and locations of the change-points. The models are described in detail in the supplementary materials, while the simulation results are provided in Table 2. EBS outperformed in many cases both in location accuracy and in controlling the overall number of change-points. What is interesting to note is that the relative performance of EBS against BS or WBS remained high in every scenario and (almost) never falls behind the other two.

6 Application

In this section, we test our proposed methodology using Apple Inc Contract For Differences (CFD) data obtained from Dukascopy for free. CFDs are leveraged products which normally reflect the underlying stock price, but they do not transfer stock ownership to the buyer. They are favoured for their easiness to access (reducing the need to trade in a stock exchange) and the extended trading hours many brokers provide. The data consists of the bid and ask quotes with their timestamps as published by the broker. Alternatively, one could choose transaction data which will require further cleansing due to the existence of multiple small size trades. We analyse all the data from June (20 business days) and from 09:30 until 16:00 (US East local time), the stock exchange trading hours. The daily average sample size (number of quotes per business day) is 60,917, the minimum 42,007, and the maximum 86,777.

For each day in month June we ran EBS (assuming a tvACD process) on the raw durations, the time lapsed between two quotes updates. This is to show whether EBS was able to segment a given series without first removing the intraday seasonality, as a piecewise constant model can approximate it well. Our methodology is designed to capture even small deviations from stationarity. Therefore, it can potentially work well in this setup whereby the underlying dynamics of a duration series change slowly and monotonically at the start of the trading day, flatten during the day and monotonically increase towards the end (resulting in the typical U-shape).

For completeness, we also ran EBS after removing intraday seasonality from the series of durations following Engle and Russell (1998). The seasonality effect is defined as a multiplicative component in the following formula

$$\tilde{X}_t = x_t \cdot \phi(\tau)$$

where $\tilde{X}_t$ are the raw durations, $\phi(\tau)$ is the intraday seasonality effect and x_t are the standardized seasonality durations. To estimate $\hat{\phi}(\tau)$ we applied a smoothing spline

Table 2 Nonstationary processes change-point detection results

Method	Model	$E(\hat{N} - N)$	$E(\hat{N} - N)$	$E[(\hat{N} - N)^2]$	HR	exec. time
BS	M1	− 0.03	0.03	0.03	0.890	0.209
	M2	− 0.11	0.11	0.13	0.870	0.162
	M3	0.97	1.13	2.35	0.674	0.014
	M4	0.67	1.87	4.93	0.266	0.031
	M5	0.21	0.21	0.31	0.923	0.037
	M6	− 6.65	7.51	76.77	0.413	0.053
	M7	− 10.68	10.70	135.70	0.247	0.048
	M8.1	− 1.14	1.48	3.70	0.592	0.073
WBS	M8.2	− 3.82	4.22	24.48	0.535	0.083
	M1	− 0.18	0.18	0.32	0.775	0.621
	M2	− 0.96	0.96	2.58	0.701	0.829
	M3	0.09	0.67	1.13	0.719	0.112
	M4	2.27	2.51	9.33	0.327	0.371
	M5	2.52	2.52	12.58	0.584	1.047
	M6	5.11	5.17	33.57	0.593	3.710
	M7	2.90	3.42	19.96	0.539	1.590
EBS	M8.1	1.50	2.58	13.86	0.478	3.503
	M8.2	6.67	8.61	119.70	0.408	2.905
	M1	0.07	0.09	0.09	0.855	0.264
	M2	− 0.47	0.67	0.47	0.817	0.275
	M3	− 0.41	1.15	2.15	0.686	0.069
	M4	2.01	2.29	7.77	0.394	0.120
	M5	0.38	0.38	0.85	0.884	0.484
	M6	3.78	4.24	26.00	0.641	1.120
	M7	− 0.43	2.69	12.59	0.554	0.854
	M8.1	− 0.71	1.55	3.91	0.552	0.791
	M8.2	− 2.83	3.37	16.51	0.617	1.017

The error measures are self-explanatory while the hit ratio HR is described in the main text. Average execution time in seconds is given in the last column. In bold, the best or the second best method based on the metric in the respective column is reported

using cross-validation to select the penalty parameter and the degrees of freedom. On the transformed durations $x_t = \tilde{X}_t / \hat{\phi}(\tau)$, we applied the EBS algorithm.

By inspecting Fig. 2 below and Figure 6 in the supplementary material, we see that EBS is able to capture the changes in the trading behaviour during the early and late sessions. Still, when removing seasonality it is very common across the days to observe change-points between 12.00 and 13.00. We do not argue that these change-points capture seasonality left out from applying a smoothing spline. In fact, we believe that financial high-frequency data are characterized by other various types of non-stationarity which cannot be ignored. This observation has implications both for trading and risk management operations. Given the rise of the electronic trading, intraday risk monitoring will need to adjust in order to

Fig. 2 Histogram of detected change-points collected over the sample period grouped by hour. On the left, EBS has been applied on the raw durations and on the right the durations has been transformed such that to remove the intraday seasonality

incorporate previous change-points and provide more reliable forecasts in anticipation of non-stationarity.

7 Conclusion

The work has addressed the problem of detecting the change-points in the structure of an irregularly spaced time series. We proposed a new ensemble-type BS methodology, whereby we combined the change-points detected by applying BS on randomized sub-samples of the data. Doing so, we were able to detect change-points that occur frequently and in short distance between them as well as we controlled the total number of change-points that appeared to be spurious. From the simulation study in Sect. 5 it appears that the EBS mechanism performs well at this task. We believe that, apart from the good performance, EBS is a good starting point of studying other ensemble-type change-point detection methods.

Supplementary Information The online version supplementary material available at <https://doi.org/10.1007/s42952-021-00120-w>.

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

References

- Anastasiou, A., & Fryzlewicz, P. (2019). Detecting multiple generalized change-points by isolating single ones. [arXiv:1901.10852](https://arxiv.org/abs/1901.10852) (arXiv preprint).
- Baranowski, R., Chen, Y., & Fryzlewicz, P. (2019). Narrowest-over-threshold detection of multiple change-points and change-point-like features. *Journal of the Royal Statistical Society Series B*, 20, 649–672.
- Cho, H., & Fryzlewicz, P. (2012). Multiscale and multilevel technique for consistent segmentation of nonstationary time series. *Statistica Sinica*, 22, 207–229.
- Cho, H., & Korkas, K. (2018). High-dimensional garch process segmentation with an application to value-at-risk. [arXiv:1706.01155](https://arxiv.org/abs/1706.01155) (arXiv preprint).
- Davidson, J. (1994). *Stochastic limit theory: An introduction for econometricians*. OUP.
- Davis, R., Huang, D., & Yao, Y. (1995). Testing for a change in the parameter values and order of an autoregressive model. *Annals of Statistics*, 23, 282–304.
- Douc, R., Moulines, E., & Stoffer, D. (2014). *Nonlinear time series: Theory, methods and applications with R examples*. Chapman and Hall.
- Engle, R. F., & Russell, J. R. (1998). Autoregressive conditional duration: A new model for irregularly spaced transaction data. *Econometrica*, 20, 1127–1162.
- Franco, C., Roussignol, M., & Zakoian, J.-M. (2001). Conditional heteroskedasticity driven by hidden Markov chains. *Journal of Time Series Analysis*, 22(2), 197–220.
- Fryzlewicz, P. (2014). Wild binary segmentation for multiple change-point detection. *The Annals of Statistics*, 42(6), 2243–2281.
- Fryzlewicz, P. (2020). Detecting possibly frequent change-points: Wild binary segmentation 2 and steepest-drop model selection. *Journal of the Korean Statistical Society*, 20, 1–44.
- Fryzlewicz, P., & Subba Rao, S. (2011). Mixing properties of arch and time-varying arch processes. *Bernoulli*, 17(1), 320–346.
- Fryzlewicz, P., & Subba Rao, S. (2014). Multiple-change-point detection for auto-regressive conditional heteroscedastic processes. *Journal of the Royal Statistical Society Series B*, 20, 903–924.
- Harchaoui, Z., & Lévy-Leduc, C. (2010). Multiple change-point estimation with a total variation penalty. *Journal of the American Statistical Association*, 105(492), 1480–1493.
- Hawkes, A. G. (1971). Point spectra of some mutually exciting point processes. *Journal of the Royal Statistical Society Series B (Methodological)*, 33(3), 438–443.
- Inclan, C., & Tiao, G. C. (1994). Use of cumulative sums of squares for retrospective detection of changes of variance. *Journal of the American Statistical Association*, 89, 913–923.
- Killick, R., Fearnhead, P., & Eckley, I. A. (2012). Optimal detection of changepoints with a linear computational cost. *Journal of the American Statistical Association*, 107(500), 1590–1598.
- Kokoszka, P., & Teyssière, G. (2002). Change-point detection in garch models: Asymptotic and bootstrap tests. Technical report, Université Catholique de Louvain.
- Korkas, K. K. (2020). *eNchange: Ensemble methods for multiple change-point detection*. R Foundation for Statistical Computing.
- Korkas, K. K., & Fryzlewicz, P. (2017). Multiple change-point detection for non-stationary time series using wild binary segmentation. *Statistica Sinica*, 27(1), 287–311.
- Kovács, S., Li, H., Bühlmann, P., & Munk, A. (2020). Seeded binary segmentation: A general methodology for fast and optimal change point detection. [arXiv:2002.06633](https://arxiv.org/abs/2002.06633) (arXiv preprint).
- Meinshausen, N., & Bühlmann, P. (2010). Stability selection. *Journal of the Royal Statistical Society Series B (Statistical Methodology)*, 72(4), 417–473.
- Mercurio, D., & Spokoiny, V. (2004). Statistical inference for time-inhomogeneous volatility models. *Annals of Statistics*, 32, 577–602.
- Mikosch, T., & Stărică, C. (2004). Nonstationarities in financial time series, the long-range dependence, and the igarch effects. *The Review of Economics and Statistics*, 86(1), 378–390.
- Olshen, A. B., Venkatraman, E., Lucito, R., & Wigler, M. (2004). Circular binary segmentation for the analysis of array-based dna copy number data. *Biostatistics*, 5(4), 557–572.
- Roueff, F., Von Sachs, R., & Sansonnet, L. (2016). Locally stationary hawkes processes. *Stochastic Processes and their Applications*, 126(6), 1710–1743.
- Wang, T., & Samworth, R. J. (2018). High dimensional change point estimation via sparse projection. *Journal of the Royal Statistical Society Series B (Statistical Methodology)*, 80(1), 57–83.

- Xin, L., & Zhu, M. (2012). Stochastic stepwise ensembles for variable selection. *Journal of Computational and Graphical Statistics*, 21(2), 275–294.
- Yao, Y. (1988). Estimating the number of change-points via Schwarz' criterion. *Statistics and Probability Letters*, 6, 181–189.
- Zhu, M. (2015). Use of majority votes in statistical learning. *Wiley Interdisciplinary Reviews Computational Statistics*, 7(6), 357–371.

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