



HAL
open science

A new segmentation method for the homogenisation of GNSS-derived IWV time-series

Annarosa Quarello, Olivier Bock, Emilie Lebarbier

► To cite this version:

Annarosa Quarello, Olivier Bock, Emilie Lebarbier. A new segmentation method for the homogenisation of GNSS-derived IWV time-series. *Journal of the Royal Statistical Society: Series C Applied Statistics*, In press, 10.48550/arXiv.2005.04683 . hal-02957051

HAL Id: hal-02957051

<https://hal.science/hal-02957051v1>

Submitted on 4 Oct 2020

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

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

A NEW SEGMENTATION METHOD FOR THE HOMOGENISATION OF GNSS-DERIVED IWV TIME-SERIES

A. QUARELLO

O. BOCK

*Université de Paris, Institut de physique du globe de Paris, CNRS, IGN, F-75005 Paris, France.
ENSG-Géomatique, IGN, F-77455 Marne-la-Vallée, France.*

E. LEBARBIER

Laboratoire Modal'X, UPL, Univ. Paris Nanterre, France.

ABSTRACT. Homogenization is an important and crucial step to improve the usage of observational data for climate analysis. This work is motivated by the analysis of long series of GNSS Integrated Water Vapour (IWV) data which have not yet been used in this context. This paper proposes a novel segmentation method that integrates a periodic bias and a heterogeneous, monthly varying, variance. The method consists in estimating first the variance using a robust estimator and then estimating the segmentation and periodic bias iteratively. This strategy allows for the use of the dynamic programming algorithm that remains the most efficient exact algorithm to estimate the change-point positions. The statistical performance of the method is assessed through numerical experiments. An application to a real data set of 120 global GNSS stations is presented. The method is implemented in the R package GNSSseg that will be available on the CRAN.

keywords: Change-point detection; Dynamic programming; Homogenization climate series; GNSS IWV series

1. INTRODUCTION

Long records of observational data are essential to monitoring climate change and understanding the underlying climate processes. However, long time series are often affected by inhomogeneities due to changes in instrumentation, in station location, in observation and processing methods, and/or in the measurement conditions around the station (Jones *et al.*, 1986). Inhomogeneities most often take the form of abrupt changes which are detrimental to estimating trends and multi-scale climate variability (Easterling and Peterson, 1995). Various homogenization methods have been developed for the detection and correction of such change-points in the context of climate data analysis, e.g. Peterson *et al.* (1998); Caussinus and Mestre (2004); Menne and Williams (2005); Szentimrey (2008); Reeves *et al.* (2007); Costa and Soares (2009); Venema *et al.* (2012). In this paper, we are interested in ground-based Global Navigation Satellite System (GNSS) integrated water vapour (IWV) measurements. GNSS measurements provide among the most accurate and continuous IWV measurements, in all weather conditions, and have not yet been used much for climate analysis (Bevis *et al.*, 1992; Bock *et al.*, 2013; Van Malderen *et al.*, 2014; Ning *et al.*, 2016b).

In order to remove the climate signal and reveal the inhomogeneities in the GNSS measurements, it has been a common approach to compare the candidate series with a well correlated reference series. The reference series can be taken from nearby stations (i.e. observing a similar climate signal) as proposed by, e.g., Caussinus and Mestre (2004), Menne and Williams (2005), or Szentimrey (2008), or from climate model data Ning *et al.* (2016a); Bock *et al.* (2018). Since the number of stations from the International GNSS Service (IGS) is limited to a hundred or so, the construction of references series from neighboring stations is hard. The second approach is here considered using the European Center for Medium Range Forecasts (ECMWF) reanalysis ERA-Interim (Dee *et al.*, 2011) as a reference. Figure 1(a) shows an example of daily IWV data from GNSS measurements and from the ERA-Interim (ERA-I) reanalysis at station CCJM. The daily IWV data exhibit a marked seasonal variation, with values varying from 10 kgm^{-2} to 60 kgm^{-2} between winter and summer, as well as a strong day-to-day variability looking as superposed noise. When the ERA-Interim data are subtracted from the GNSS data, one clear jump can be seen on 24 Feb 2001 (Figure 1(b)). This jump coincides with a change of receiver and antenna at this station.

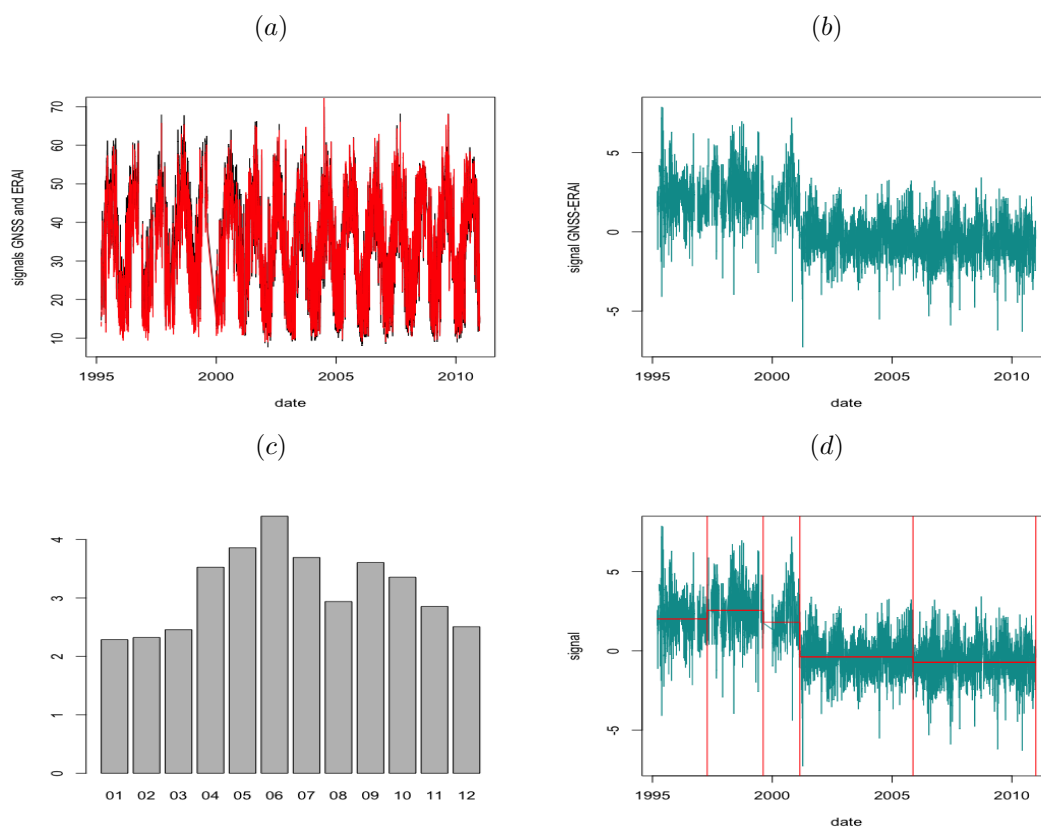


FIGURE 1. Station CCJM: (a) GNSS (in black) and ERA-Interim (in red) IWV time series; (b) IWV difference (GPS - ERA-Interim) series; (c) estimated monthly variance; (d) obtained change-points with the method proposed by Bock *et al.* (2018).

In a previous work, Bock *et al.* (2018) proposed a first segmentation method to detect abrupt changes in the mean in such data (GNSS - ERA-I IWV differences). A specific feature of their method is that it accounts for a heterogeneous variance that is assumed to vary on a monthly basis. Indeed, as can be seen in Figure 1(c), the GNSS - ERA-I IWV differences show a seasonal variation

with an increased variability in summer. Thus classical segmentation models with homogeneous or segment-specific variance are not adapted. The result of their model is given in Figure 1(d). The previously mentioned jump is well detected. However, as already mentioned in Bock *et al.* (2018), despite the ERAI data are subtracted, it can happen that not all the climate signal is removed due to representativeness differences between the reanalysis and the GNSS observations (Bock and Parracho, 2019). This residual signal exhibits a strong seasonal variation which can lead to wrong, misplaced, or missing change-points.

This paper described an improved method which accounts for seasonal variation in the signal by adding a functional term to earlier model used by Bock *et al.* (2018). To infer the parameters of this enhanced model, a (penalized)-maximum likelihood procedure is used again. In this framework, it is well known that segmentation methods have to deal with two problems: (i) an inherent algorithmic complexity for estimating the change-point locations and (ii) an appropriate choice of the penalty term which controls the number of change-points. Indeed, for problem (i), the inference of the discrete change-points requires to search over the whole segmentation space that is huge. Such a search is prohibitive in terms of computational time when performed in a naive way. The Dynamic Programming (DP) algorithm (Auger and Lawrence, 1989) and its recent pruned versions (Killick *et al.*, 2012; Rigail, 2015; Maidstone *et al.*, 2017), are the only algorithms that retrieve the exact solution in a fast way. However, a necessary condition for using DP is that the quantity to be optimized is additive with respect to the segments (Bai and Perron, 2003; Caussinus and Mestre, 2004; Picard *et al.*, 2005). Here, with the presence of the monthly variance and the functional part, the condition is not verified. To circumvent this, Li and Lund (2012) and Lu *et al.* (2010) proposed to use a genetic algorithm. However, this algorithm leads to a suboptimal solution.

Our objective here is to keep the interest of the exact DP algorithm as possible. To be able us to use it in the inference procedure, we propose to (1) estimate first the variance using a robust (to the change-points) estimator as in Bock *et al.* (2018) and (2) treat sequentially the estimation of the segmentation parameters and the functional as in Bertin *et al.* (2017). For the choice of the number of segments, different penalties have been proposed in the literature (see Lebarbier (2005); Lavielle (2005); Zhang and Siegmund (2007); Lu *et al.* (2010); Caussinus and Mestre (2004)). Here we propose to use some of them.

The article is organized as follows. Section 2 presents the model and the inference procedure. A simulation study is performed in Section 3 to evaluate the performance of the method. In Section 4 the method is applied on real data from a set 120 global GNSS stations. Section 5 discusses the results and concludes.

2. MODEL AND INFERENCE

2.1. Model. We consider the model proposed by Bock *et al.* (2018) in which we add a functional part in order to take into account the periodic bias. Let $\mathbf{y} = \{y_t\}_{1,\dots,n}$ be the observed series with length n that is modeled by a Gaussian independent random process $\mathbf{Y} = \{Y_t\}_{t=1,\dots,n}$ such that

- the mean of \mathbf{Y} is composed of two terms:
 - ★ a piecewise constant function equals to μ_k on the interval $I_k^{\text{mean}} = \llbracket t_{k-1} + 1, t_k \rrbracket$ with length $n_k = t_k - t_{k-1}$ where $0 = t_0 < t_1 < \dots < t_{K-1} < t_K = n$. The $\{t_k\}_{k=1,\dots,K-1}$ are the times of the change-points and K is the number of intervals or segments,
 - ★ and a function f ;
- the variance of \mathbf{Y} is month-dependent, i.e. it is constant within the interval $I_{\text{month}}^{\text{var}} = \{t; \text{date}(t) \in \text{month}\}$ with length n_{month} where $\text{date}(t)$ stands for the date at the time t .

The resulting model is thus the following

$$(1) \quad Y_t = \mu_k + f_t + E_t, \quad \text{where } E_t \sim \mathcal{N}(0, \sigma_{\text{month}}^2) \text{ if } t \in I_k^{\text{mean}} \cap I_{\text{month}}^{\text{var}},$$

for $k = 1, \dots, K$. The intervals $\{I_k^{\text{mean}}\}_k$ are unknown contrary to the intervals $\{I_{\text{month}}^{\text{var}}\}_{\text{month}}$. The parameters to be estimated are the number of segments K (or the number of change-points $K - 1$), the change-points $\mathbf{t} = \{t_k\}_k$ and the distribution parameters, the means $\boldsymbol{\mu} = \{\mu_k\}_k$, the variances $\sigma^2 = \{\sigma_{\text{month}}^2\}_{\text{month}}$ and the function f .

2.2. Inference. As usual in segmentation, the inference is performed in two steps (e.g. Truong *et al.* (2020)):

Step 1: Estimate \mathbf{t} , $\boldsymbol{\mu}$, σ^2 and f , K being fixed.

Step 2: Choose the number of segments K .

We consider here a penalized maximum likelihood approach. The log-likelihood of the model defined by (1) is

$$(2) \quad \log p(\mathbf{y}; K, \mathbf{t}, \boldsymbol{\mu}, \sigma^2, f) = -\frac{n}{2} \log(2\pi) - \sum_{\text{month}} \frac{n_{\text{month}}}{2} \log(\sigma_{\text{month}}^2) - \frac{1}{2} \sum_{k=1}^K \sum_{\text{month}} \sum_{t \in I_k^{\text{mean}} \cap I_k^{\text{var}}} \frac{(y_t - \mu_k - f_t)^2}{\sigma_{\text{month}}^2}$$

2.2.1. Step 1: Inference of \mathbf{t} , $\boldsymbol{\mu}$, σ^2 and f , K being fixed. The use of the DP algorithm is now classical to estimate the change-points. However, DP can be applied if and only if the quantity to be optimized is additive with respect to the segments. Here the presence of the 'global' parameters σ_{month}^2 and f will link the segments and the required condition will not be satisfied. To circumvent this problem a two-step procedure is proposed: (1) we estimate the variances using a robust estimator as in Chakar *et al.* (2017) and Bock *et al.* (2018) and (2) we estimate iteratively f and the segmentation parameters (i.e. the change-points and the means) using DP as in Gazeaux *et al.* (2015) and Bertin *et al.* (2017).

The resulting algorithm is the following:

Estimation of σ_{month}^2 : Bock *et al.* (2018) proposed a consistent estimator for the variance parameter based on the robust one proposed by Rousseeuw and Croux (1993). The key idea is to apply this robust estimator (up to a constant) on the differentiated series $y_t - y_{t-1}$. This series is centered except at the change-point positions which are treated as outliers. We again use this estimator even in the presence of the function f because the latter does not have much impact on the resulting estimation (in the application, the seasonal signal is slowly varying and is almost completely cancelled in the differentiated series). The estimated variance is noted $\hat{\sigma}_{\text{month}}^2$.

Estimation of f and both \mathbf{t} and $\boldsymbol{\mu}$ iteratively: by minimizing the minus log-likelihood given in (2). At iteration $[h+1]$:

- (a) The estimator of f results in a weighted least-square estimator with weights $1/\hat{\sigma}_{\text{month}}^2$ on $\{y_t - \mu_k^{[h]}\}_t$. For our application and following Weatherhead *et al.* (1998), we represent f as a Fourier series of order 4 accounting for annual, semi-annual, terannual, and quarterly periodicities in the signal:

$$f_t = \sum_{i=1}^4 a_i \cos(w_i t) + b_i \sin(w_i t),$$

where $w_i = 2\pi \frac{i}{L}$ is the angular frequency of period L/i and L is the mean length of the year ($L = 365.25$ days when time t is expressed in days). The estimated function is denoted $f^{[h+1]}$.

- (b) The segmentation parameters are estimated based on $\{y_t - f_t^{[h+1]}\}_t$. We get

$$\mu_k^{[h+1]} = \frac{\sum_{\text{month}} \sum_{t \in I_k^{\text{mean}} \cap I_k^{\text{var}}} \frac{(y_t - f_t^{[h+1]})}{\hat{\sigma}_{\text{month}}^2}}{\sum_{\text{month}} \sum_{t \in I_k^{\text{mean}} \cap I_k^{\text{var}}} \frac{1}{\hat{\sigma}_{\text{month}}^2}},$$

and

$$\mathbf{t}^{[h+1]} = \underset{\mathbf{t} \in \mathcal{M}_{K,n}}{\operatorname{argmin}} \sum_{k=1}^K \sum_{\text{month}} \sum_{t \in I_k^{\text{mean}} \cap I_k^{\text{var}}} \frac{(y_t - f_t^{[h+1]} - \mu_k^{[h+1]})^2}{\hat{\sigma}_{\text{month}}^2},$$

where $\mathcal{M}_{K,n} = \{(t_1, \dots, t_{K-1}) \in \mathbb{N}^{K-1}, 0 = t_0 < t_1 < \dots, t_{K-1} < t_K = n\}$ is the set of all the possible partitions of the grid $\llbracket 1, n \rrbracket$ in K segments. This minimization is obtained using DP.

The final estimators are denoted \hat{f} , $\hat{\mathbf{t}}$ and $\hat{\boldsymbol{\mu}}$.

2.2.2. *Choice of K .* Various criteria have been theoretically developed for the choice of K in segmentation with a homogeneous (known or unknown) variance. However, no criteria exist for the case with a heterogeneous variance on fixed intervals. Since in our estimation procedure the variances are estimated first, our segmentation problem can be seen as one in which the variance is known. We thus propose to use the least-squares-based criterion:

$$(3) \quad \text{SSR}_K(\hat{\mathbf{t}}, \hat{\boldsymbol{\mu}}, \hat{\sigma}^2, \hat{f}) = \sum_{k=1}^K \sum_{\text{month}} \sum_{t \in \hat{I}_k^{\text{mean}} \cap I_{\text{month}}^{\text{var}}} \frac{(y_t - \hat{f}_t - \hat{\mu}_k)^2}{\hat{\sigma}_{\text{month}}^2}.$$

Different penalties are considered and tested in this paper:

Lav: proposed by Lavielle (2005):

$$\hat{K} = \underset{K}{\text{argmin}} \text{SSR}_K(\hat{\mathbf{t}}, \hat{\boldsymbol{\mu}}, \hat{\sigma}^2, \hat{f}) + \beta K,$$

where β is the penalty constant chosen using an adaptive method. The method involves a threshold S which is fixed to $S = 0.75$, both in the simulation study and the applications, as suggested by Lavielle (2005).

BM: proposed by Birgé and Massart (2001) and Lebarbier (2005) for an application in a segmentation context:

$$\hat{K} = \underset{K}{\text{argmin}} \text{SSR}_K(\hat{\mathbf{t}}, \hat{\boldsymbol{\mu}}, \hat{\sigma}^2, \hat{f}) + \alpha K \left[5 + 2 \log \left(\frac{n}{K} \right) \right],$$

where the penalty constant α can be calibrated using the slope heuristic proposed by Arlot and Massart (2009). Two methods are proposed actually: the "dimension jump" and the "data-driven slope estimation" which are referred to as BM1 and BM2, respectively, hereafter.

mBIC: the modified version of the classical BIC criterion derived in the segmentation framework by Zhang and Siegmund (2007),

$$\hat{K} = \underset{K}{\text{argmax}} -\frac{1}{2} \text{SSR}_K(\hat{\mathbf{t}}, \hat{\boldsymbol{\mu}}, \hat{\sigma}^2, \hat{f}) - \frac{1}{2} \sum_{k=1}^K \log(\hat{t}_k - \hat{t}_{k-1}) + \left(\frac{1}{2} - K \right) \log(n).$$

In the specific climate context, some authors as Li and Lund (2012) and Lu *et al.* (2010) use a MDL based-criterion (Rissanen (1978)). Ardia *et al.* (2019) show that the MDL criterion can be seen as a Bayesian criterion with appropriate prior distributions for change-point models. As a consequence, the obtained based-MDL penalties (see Li and Lund (2012); Lu *et al.* (2010)) looks like the mBIC (their both penalties integrate a term depending on the segment lengths of the segmentation).

2.2.3. *Different choices for our procedure.* The proposed inference procedure is summarized in Figure 10 given in the Supplemental Material. The method is implemented in a R package named `GNSSseg` which is available on the CRAN.

In practice, Step 1 of the inference (Section 2.2.1) is performed for $K = 1, \dots, K_{\max}$ where K_{\max} should be 2 or 3 times larger than the expected number of change-points. For both the simulations and the applications, we used $K_{\max} = 30$.

The iterative procedure needs a proper initialization procedure and a stopping rule. For the initialization, the function f is estimated first, using a unweighted least-squares criterion. For the stopping rule the change of f_t and μ_k between two successive iterations is checked against a fixed threshold. The convergence of the iterative procedure is accelerated using the stopping test proposed by Varadhan and Roland (2008).

The final parameterization was derived after testing several different options which are described in the Supplemental Material.

3. SIMULATION STUDY

3.1. Simulation design and quality criteria.

Simulation design. The simulated time series are characterized by a length of $n = 400$ with 4 "years" of 2 "months" of 50 "days" each and with a monthly variance. A total of 6 change-points were introduced at positions $t = 55, 77, 177, 222, 300, 366$ and values for the signal mean were alternating between 0 and 1. The periodic function was modelled by $f_t = 0.7 \cos(2\pi t/L)$ where $L = 100$ is the length of one year. Since we consider here only two months, the variance is alternating between two values, σ_1^2 and σ_2^2 . Several batches of 100 time series were generated with different values for $\sigma_1 = 0.1, 0.5, \text{ or } 0.9$ and $\sigma_2 = 0.1$ to 1.5 by step of 0.2. Figure 2 shows an example.

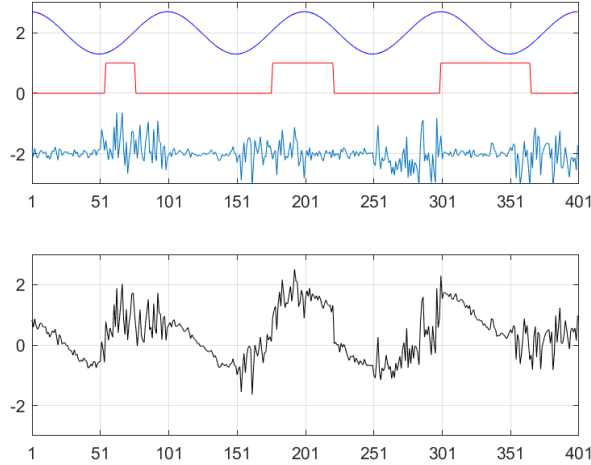


FIGURE 2. Example of a simulated time series (black solid line in lower panel) of length $n = 400$ with $K = 7$ segments (red solid line), function $f_t = 0.7 \cos(2\pi t/L)$ (blue solid line), noise (cyan solid line) with standard deviation $\sigma_1^* = 0.1$ and $\sigma_2^* = 0.5$ (changing every $L/2 = 50$ points, starting with σ_1^*).

Quality criteria. The accuracy of the results is quantified by the differences between the estimates (denoted with a hat \hat{x}) and the true values (denoted as x^*).

For the function f , the root mean square error (RMSE) of the estimated function is computed:

$$\text{RMSE}(f) = \left[\frac{1}{n} \sum_{t=1}^n \{\hat{f}_t - f_t^*\}^2 \right]^{1/2}.$$

For the segmentation parameters, the following criteria are considered:

- the difference between the estimated number of segments and the true one $\hat{K} - K^*$;
- the RMSE of the estimated mean parameter $\hat{\mu}$: $\text{RMSE}(\mu) = \left[\frac{1}{n} \sum_{t=1}^n \{\hat{\mu}_t - \mu_t^*\}^2 \right]^{1/2}$;
- the distance between the estimated positions of the change-points $\hat{\mathbf{t}}$ and the true ones \mathbf{t}^* ; this distance is measured with the help of the two components of the Hausdorff distance, d_1 and d_2 , defined as:

$$d_1(a, b) = \max_b \min_a |a - b| \quad \text{and} \quad d_2(a, b) = d_1(b, a).$$

A perfect segmentation results in both null $d_1(\mathbf{t}^*, \hat{\mathbf{t}})$ and $d_2(\mathbf{t}^*, \hat{\mathbf{t}})$. A small d_1 means that the detected change-points are well positioned and a small d_2 that a large part of the true change-points are correctly detected. A common situation found in practice is the one where the number of change-points is under-estimated, with a small d_1 and a large d_2 . In that case, some change-points are undetected but the detected ones are correctly located. This situation is satisfying here since in our application it is acceptable to miss a

few change-points (usually of small amplitude) rather than over-segmenting the data with badly-positioned change-points.

- the histogram of the change-point locations that provides a measure of the probability of the position of the change-points.

3.2. Results. Only the results for $\sigma_1^* = 0.5$ are illustrated hereafter. The results for the others values of σ_1^* are briefly discussed at the end of the section.

Accuracy of the variance estimates. Figure 3 presents the estimation errors of $\hat{\sigma}_1$ and $\hat{\sigma}_2$ for different values of σ_2^* . It is seen that the variance estimator works well and the estimated standard deviations are retrieved with the same accuracy as in Bock *et al.* (2018) despite the presence of the periodic bias. The dispersion increases when σ_2^* is increasing as one can expect.

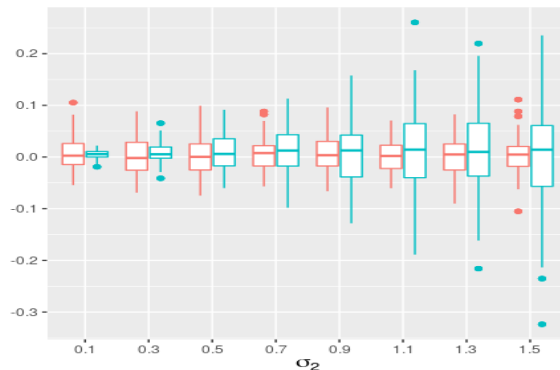


FIGURE 3. Boxplots of standard deviation estimation errors: $\hat{\sigma}_1 - \sigma_1^*$ in red and $\hat{\sigma}_2 - \sigma_2^*$ in blue, with $\sigma_1^* = 0.5$ and $\sigma_2^* = 0.1, \dots, 1.5$. Each case includes 100 simulations.

Accuracy of segmentation parameter estimates. Figure 4 shows the results for the four model selection criteria and the special case where the number of segments K is fixed to the true value ($K = 7$). For small values of σ_2^* , the detection problem is easy and all the model selection criteria retrieve the correct number of segments (Figure 4(a)). However for large values of σ_2^* , the detection becomes difficult, and the errors increase. The different criteria behave slightly differently. Lav tends to give the true number of segments in median, but with a large dispersion, while BM1, BM2, and mBIC tend to underestimate the number of segments (more for mBIC). However, finding the correct number of segments does not mean that the change-points are properly positioned. Indeed, for Lav and the case when $K = 7$, the median d_1 is still quite large (Figure 4(c)). On the other hand, the median d_2 is smaller for the case when $K = 7$ compared to the tested criteria (Figure 4(d)). Finally, RMSE(μ) is very similar for all the criteria (Figure 4(b)), though Lav shows a larger median and dispersion when σ_2^* is large. When σ_2^* takes intermediate values the case when $K = 7$ yields slightly improved results.

Probability of detection. Figure 5 shows the percentage of the change-point detections for three values of $\sigma_2^* = 0.1, 0.5$ and 1.5 , and $\sigma_1^* = 0.5$. In general, the change-points located in the "months" with smaller variance are more often recovered with all three criteria, and also when the true K is used. Hence, in the case (a) when $\sigma_1^* = 0.5$ and $\sigma_2^* = 0.1$, the probability of detection is slightly smaller for the position 222, which is contained in a segment with $\sigma_1^* = 0.5$, and for the position 300 where both the mean and the variance change. In the case (b) when $\sigma_1^* = \sigma_2^* = 0.5$, the probability of detection is more or less the same for all the change-points and all the criteria. When $\sigma_2^* = 1.5$, the problem is more complicated. Again the change-points located in the "months" with smaller noise are better detected (positions 222 and 300) but for the other four change-points the results are contrasted although they are all located in months with $\sigma_2^* = 1.5$. The change-points at 55 and 77 are almost never detected. For mBIC this is consistent with the fact that the median $\hat{K} = 5$, i.e.

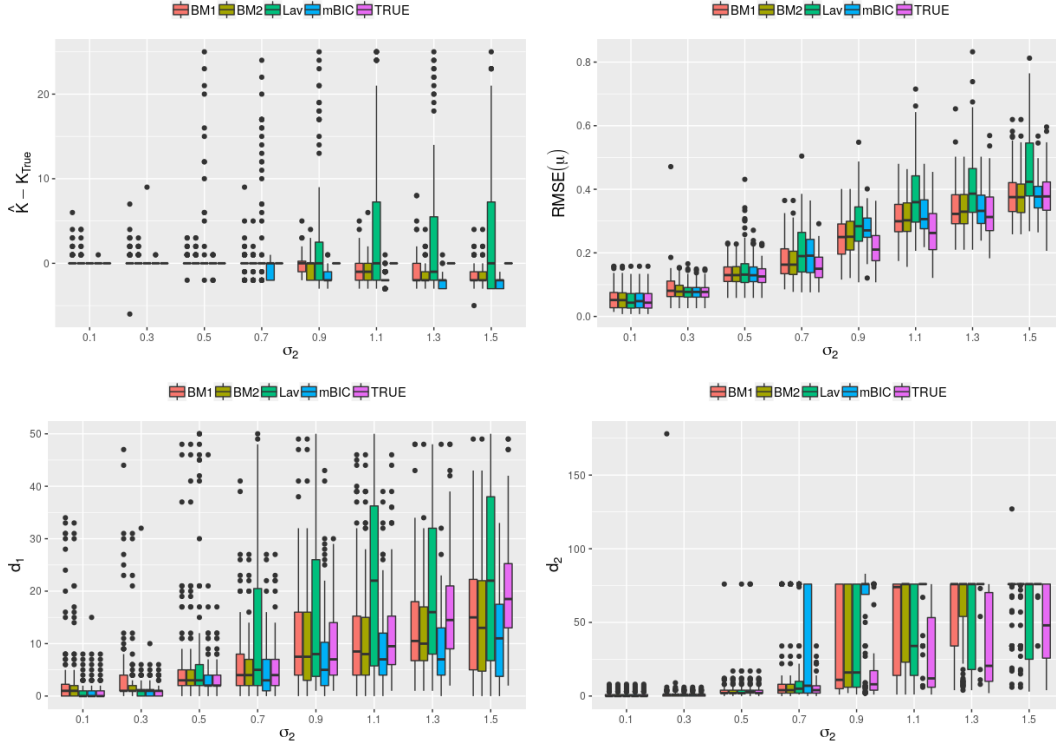


FIGURE 4. Results with the four selection criteria (BM1, BM2, Lav, and mBIC) and with the true number of segments (True), for $\sigma_1^* = 0.5$ and different values of σ_2^* . (a) $\hat{K} - K^*$; (b) $\text{RMSE}(\mu)$; (c) first Hausdorff distance d_1 and (d) second Hausdorff distance d_2 .

two change-points are missing, on average (Figure 4(a)), but the other four change-points are not so badly located (d_1 is not that large, Figure 4(c), but d_2 is very large, Figure 4(d)). The situation is a bit similar for BM1. On the other hand, for Lav and the true K , the number of detections is correct (on average for Lav) but due to the large noise they are sometimes very badly positioned (large d_1 and d_2).

Accuracy of the function estimate. Figure 6 shows $\text{RMSE}(f)$ as a function of σ_2^* . As expected, the errors increase when σ_2^* increases. The results do not much depend on the selection criterion, but the results are slightly better when the true number of segments is known and when σ_2^* takes intermediate values. The results for Lav show a slightly larger median and larger dispersion.

The results for other values of σ_1^* (not shown) are very similar for BM1, BM2, mBIC, and the case when the true K is used. The results are slightly improved for $\sigma_1^* = 0.1$ and slightly degraded for $\sigma_1^* = 0.9$, as expected. The results for Lav are more chaotic, with either large under-estimation of K for the smaller σ_1^* and over-estimation of K for the larger σ_1^* , with large subsequent degradation of the other quality criteria. In general, under-estimating K leads to an increase of $\text{RMSE}(\mu)$, while over-estimating K leads to an increase of d_1 .

The main conclusions from the simulation study are the following:

- The proposed method works well but the results are sensitive to the choice of the function form due to its possible confusion with the change-points. Performing a selection of the statistically significant parameters of the function appears as a good way to reduce this problem and improves slightly the change-point detection with our simulated data (see Supplemental Material).

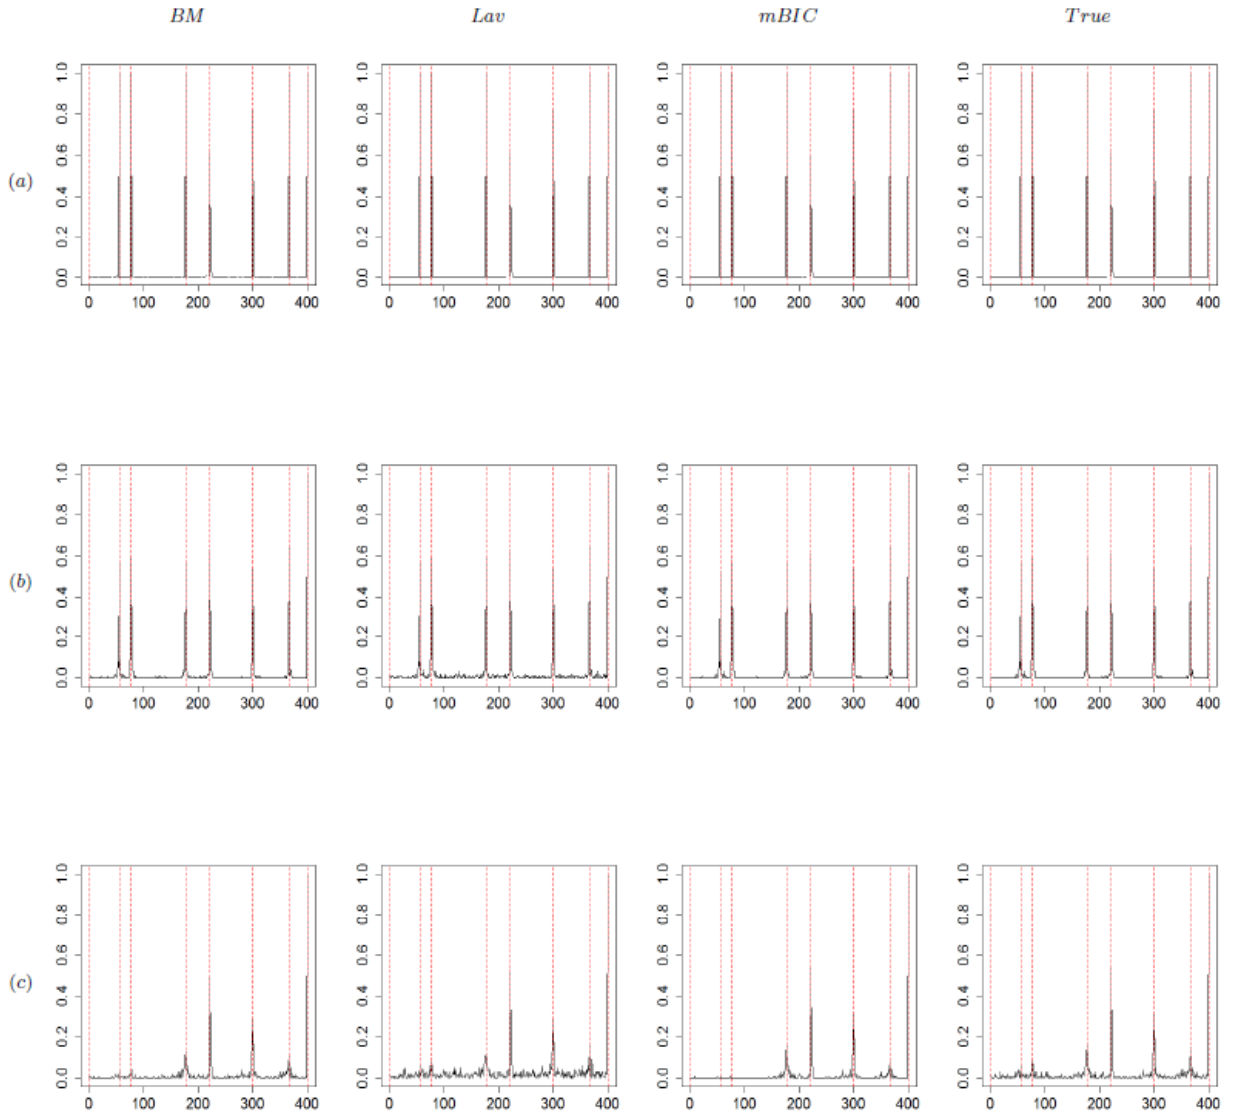


FIGURE 5. Histogram of change-point detections with, from left to right, the BM1, Lav, and mBIC selection criteria, and the case when the true number of segments is used (TRUE), for $\sigma_1^* = 0.5$ and three different values for σ_2^* : (a) $\sigma_2^* = 0.1$, (b) $\sigma_2^* = 0.5$ and (c) $\sigma_2^* = 1.5$. The red dotted lines indicate the positions of the true change-points. The results for BM2 are very similar to BM1 and are not shown.

- Concerning the model selection criteria, BM1, BM2, and mBIC, provide very similar results. They behave well and detect correctly the number and position of change-points when the noise is not too large. When the noise is heavy some change-points are missed but this is a counterpart of the limited number of false detections. The Lav criterion shows much larger dispersion in the number of change-points and, though the estimated number is close to the truth in median, some change-points are not properly located (larger d_1 and d_2) with an impact on the estimated μ and f .

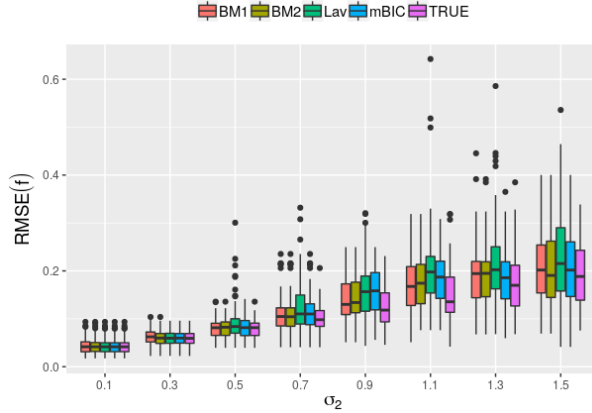


FIGURE 6. RMSE of the estimated function f for $\sigma_1^* = 0.5$ and different values for σ_2^* .

4. APPLICATION TO REAL GNSS DATA

4.1. Dataset, metadata, and validation procedure. The method is applied to the daily IWV differences from 120 global GNSS stations and ERA-Interim reanalysis for the period from 1 January 1995 to 31 December 2010. The dataset used here is the same as in Bock and Parracho (2019). The metadata for the GNSS stations are available from the IGS site-logs (<ftp://igs.org/pub/station/log/>). They contain for each station the dates of changes of receiver (R), antenna (A), and radome (D). We also included the dates of processing changes (P) which occurred at a few stations in 2008 and 2009 (this issue is discussed in Parracho *et al.* (2018)). Experience shows that equipment changes do not produce systematically a break in the GNSS IWV time series. The most important changes are those affecting the antenna and its electromagnetic environment, the satellite visibility, and the number of observations (Vey *et al.*, 2009). For instance, Ning *et al.* (2016a) considered only antenna and radome changes, as well addition/removal of microwave absorbing material which was known by the authors for one specific station. However, there is some evidence that changes in the receiver settings also induce inhomogeneities, e.g. when the elevation cutoff angle is changed. Changes in the environment due e.g. to cutting of vegetation and construction of buildings nearby the antenna as well as seasonal changes in multipath due to growing/declining vegetation may also impact the measurements and produce either abrupt or gradual changes. As a consequence, though metadata represent a valuable source of validation, a full matching between detected change-points and metadata is not to be expected.

Because of noise in the signal, the detected changes may also not coincide perfectly with the known changes and we must allow some flexibility in the validation procedure. A window of 30 days before or after a documented change was used for the automatic validation of the detected change-points. A visual inspection was also performed to check if the invalidated change-points make sense. In some cases double detections just a few days apart are found on noise spikes, often with two large offsets of opposite signs. Such noise detections are classified as outliers.

4.2. General Results. In this section, we present results for the final method described in the preceding sections as well as for three alternative methods. The final method is referred to as variant (a). Variant (b) is a similar method where only the statistically significant terms of the Fourier series are selected. It is intended to check if reducing the number of degrees of freedom in the function leads to better results as was found with the simulations. Variant (c) is the earlier method proposed by Bock *et al.* (2018) in which only the segmentation is performed (i.e. the functional part is removed). Variant (d) is another form of a simplified method where the functional is modelled but a homogeneous variance is considered instead of a monthly variance. Statistics on the number of detected change-points are included in Figure 7. More statistics including the number of validations and outliers are given in Table 1.

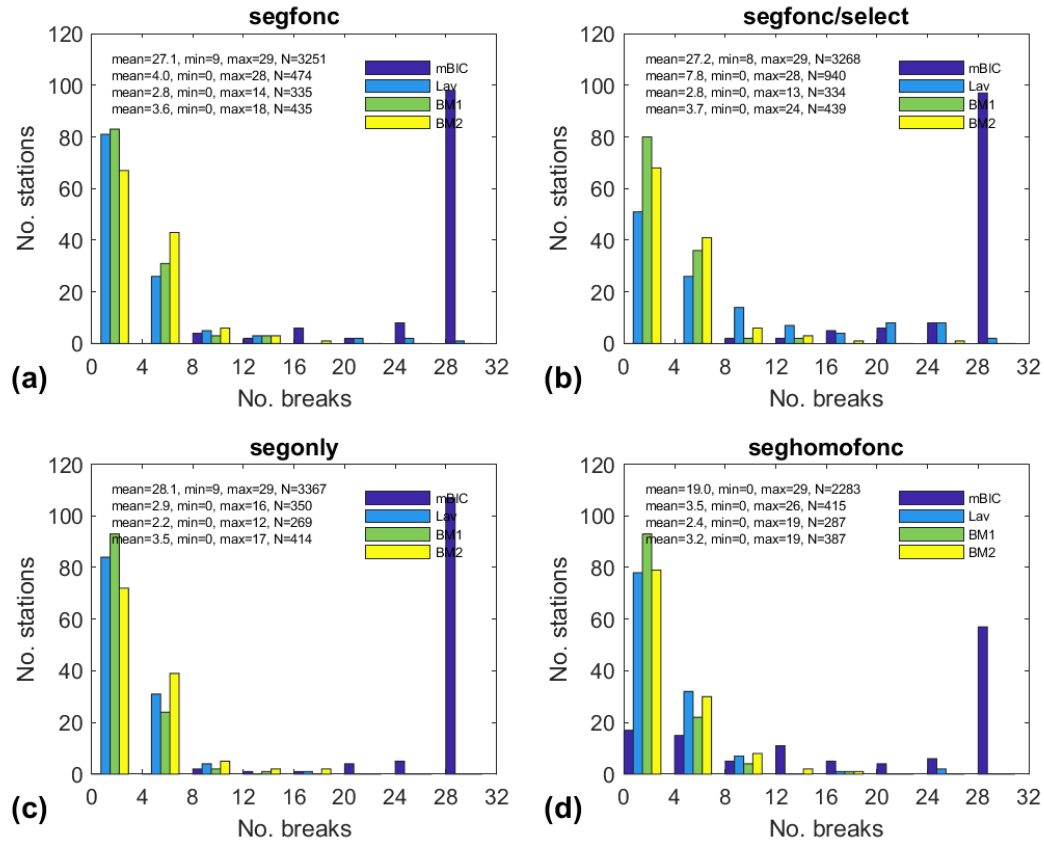


FIGURE 7. Histograms of the number of change-points detected for four variants of the model selecting criteria (mBIC, Lav, BM1, and BM2). The numbers given in the plots are the mean, min, and max number of change-points detected per station, N is the total number of change-points per method.

Figure 7(a) shows that with variant (a), mBIC, Lav, BM1, and BM2 detect a total of 3251, 474, 335, and 435 change-points, respectively. The distribution of the number of change-points per station is very different depending on the selection criterion. Most notably, mBIC detects between 9 and 29 change-points per station, with a mean value of 27.1, i.e. in most cases the maximum number of segments is selected (here $K_{\max} = 30$). This behaviour was not observed with the simulations. From Table 1 we see that mBIC detects many outliers. Comparison of contrast values reveals that mBIC selects solutions with smaller SSR values than the other criteria, i.e. the model selected by mBIC generally explains better the observed signal. However, this is at the expense of strong over-segmentation, which is not wanted. mBIC is thus not well adapted to the nature of the data analyzed here. One of the reasons might be that the hypothesis of Gaussian errors is not valid (e.g. due to serial correlation in the data and noise spikes). The three other selection criteria provide much more consistent results, with mean number of change-points of 2.8, 3.6 and 4.0 for BM1, BM2, and Lav, respectively. Among the three criteria, we see from Table 1 that BM1 has the smallest number of outliers (36) and the highest rate of validations (20.9%). These two features, and also the fact that BM1 has a reasonable number of change-points (the mean is 2.8 per station), make this selection criterion the preferred one.

Compared to variant (a), variant (b) shows marginal impact on the number of detections and the number of validations for three criteria (mBIC, BM1, and BM2). Only for Lav do the mean and total number of detections increase (by nearly a factor of 2). This behavior is not explained

TABLE 1. Comparison of segmentation results for the four variants and the four model selection criteria. From left to right: Number of stations with change-points, min/mean/max number of detected change-points per station, total number of change-points, total number of outliers, total number of validations, percentage of validations including outliers, percentage of validations without outliers.

	Nsta	min	mean	max	detections	outliers	validations		
Variant (a) (segfonc)									
<i>mBIC</i>	120	9	27.1	29	3251	2096	267	8.2%	20.9%
<i>Lav</i>	114	0	4.0	28	474	129	75	15.8%	21.3%
<i>BM1</i>	98	0	2.8	14	335	36	70	20.9%	23.3%
<i>BM2</i>	107	0	3.6	18	435	64	77	17.7%	20.6%
Variant (b) (segfonc/select)									
<i>mBIC</i>	120	8	27.2	29	3268	2090	270	8.3%	20.7%
<i>Lav</i>	115	0	7.8	28	940	411	116	12.3%	20.8%
<i>BM1</i>	100	0	2.8	13	334	46	68	20.4%	23.4%
<i>BM2</i>	107	0	3.7	24	439	76	81	18.5%	22.1%
Variant (c) (segonly)									
<i>mBIC</i>	120	9	28.1	29	3367	1255	361	10.7%	16.4%
<i>Lav</i>	113	0	2.9	16	350	28	64	18.3%	19.6%
<i>BM1</i>	90	0	2.2	12	269	8	53	19.7%	20.2%
<i>BM2</i>	102	0	3.5	17	414	24	68	16.4%	17.4%
Variant (d) (seghomofonc)									
<i>mBIC</i>	116	0	19.0	29	2283	1637	178	7.8%	24.1%
<i>Lav</i>	114	0	3.5	26	415	148	56	13.5%	20.4%
<i>BM1</i>	92	0	2.4	19	287	40	61	21.3%	24.1%
<i>BM2</i>	101	0	3.2	19	387	82	68	17.6%	21.7%

but it reveals some instability in the model selection with this criterion. Instability could also be guessed from the maximal number of detections of 28 already seen in variant (a). It means that in some cases, *Lav* selects a number of segments very close to the maximum ($K_{\max} = 30$). *BM1* and *BM2* have also more outliers with this variant, though the total number of detections is almost unchanged. So, contrary to the simulation results, with the real data there is no benefit of applying a selection of significant terms of the functional model.

In variant (c), the result for *mBIC* is slightly worse (more detections) but with fewer outliers. For the three other criteria the number of detections decreases significantly. The latter behaviour was actually not expected. Our interpretation is that when the periodic bias is not modelled, the segmentation algorithm has two options: either (i) put additional change-points to better fit the periodic variations in the signal, but this would lead to many more detections (4 per year, i.e. a total of 64 per station for a 16-year time series), or (ii) select only those change-points with a large amplitude that are not confounded with the periodic bias. The observed result (Figure 7(c) and Table 1) suggest that *BM1*, *BM2*, and *Lav* select the second, more conservative, option. Our final method is actually capable of detecting smaller offsets, which makes it more efficient for the homogenization purpose. Note that with variant (c), the situation described by option (i) occurs nevertheless in some cases, as will be illustrated in the next sub-section, and though the number of outliers and validations both decrease for *BM1*, *BM2*, and *Lav*, the percentage of validations remains nearly the same (Table 1). So, variant (a) clearly works better than variant (c) in the sense it detects more change-points; it has nevertheless the drawback of detecting more outliers. This point is further discussed in the last section.

In variant (d) the variance is assumed to be constant. This has two consequences: (i) the function is fitted with uniform weights which in general leads to an estimated function \hat{f} and an estimated mean $\hat{\mu}$ of different shapes, (ii) the estimated variance is larger than the mean variance of the variant (a) (the average mean standard deviations amount to 1.19 vs. 0.84kgm^{-2} , respectively) and fewer change-points are detected. Table 1 confirms that with this method fewer change-points are detected than with variant (a), however the number of outliers is increased (except for mBIC which is again a special case). The number of validations is also decreased, but the percentage of validations is almost unchanged.

The comparison of the four variants shows thus that the final method, including a heterogeneous variance and a full functional model for the periodic bias, has the best properties: reasonable number of detections and outliers, and high rate of validations. Among the four model selection criteria, BM1 and BM2 behave better than Lav and mBIC, with a small advantage for BM1. Figure 8(a) shows that the yearly-mean standard deviation of the noise ranges between 0 and 2kgm^{-2} , with a mean value over the 120 stations of 0.84kgm^{-2} . The seasonal excursion is of 0.63kgm^{-2} on average, which reflects the importance of modelling the heterogeneous variance. Figure 8(b) presents a measure of the magnitude of the periodic bias for BM1. With an average value of 0.33kgm^{-2} it is clear that the periodic bias is not negligible and modelling it improves the segmentation results as shown by comparing the results of variant (d) and (a). Figure 8(c) shows that the distribution of offsets (changes in mean) is nearly symmetrical. The mean absolute value of 1.27kgm^{-2} is relatively large. The dip centred on zero reflects the fact that the smaller offsets are more difficult to detect because of their small signal-to-noise (SNR) ratio. The most frequently detected offsets are found around $\pm 0.5\text{kgm}^{-2}$. The larger offsets (up to $\pm 10\text{kgm}^{-2}$) are outliers. The distribution of SNR can be computed as the absolute value of offset divided by standard deviation of noise. It is peaking at 0.6 and the larger values (up to 10) correspond again to outliers (Figure 8(d)). The mean SNR of 1.55 indicates that our method has a good efficiency of detection.

4.3. Examples of special cases. In addition to the global results, we exhibit the results for four stations showing for special cases of the variants. Only the criterion BM1 is considered here. With variant (c) there are actually 66 stations which have the same number of detections as variant (a). Though in general the change-points are located at the same position in the time series, this is not always the case. For 18 stations, variant (c) detects more change-points and for 36 stations it detects fewer. Station POL2 is an example of the former category and station STJO an example of the latter. DUBO is an example where the same number is detected but the change-points are not located at the same position. With variant (d), the number of stations with equal, more, and fewer numbers of detections is: 57, 24, and 39, respectively. Examples are: EBRE, MCM4, and POL2, respectively.

The results for a selection of four stations are given in Figure 9:

- In the case of POL2, variants (a), (c) and (d) detect 3, 12, and 1 change-point, respectively. The signal shows a strong periodic variation which well fitted by the models of variant (a) and (d) but is erroneously captured by the segmentation in variant (c). Variant (a) has one validated change-point (detected date: 2008-02-23, known change: 2008-03-06, type of change: P). Variant (c) has no validation, although it detects 12 change-points. Variant (d) detects only one change-point, which is located 72 days from the nearest known change-point and is thus not validated, but it coincides with one of the three detections found by variant (a). The detection of this change-point is made difficult because it is located in a month with heavy noise.
- In the case of STJO, variants (a) and (d) detect 5 and 4 change-points, respectively, with one outlier each but not at the same position. Among the detected change-points, one is exactly the same (detected: 2003-04-18, known: 2003-06-08, type: R) but is not validated, and one is close (detected by variant (a): 1999-07-20, by variant (d): 1999-07-19, known:

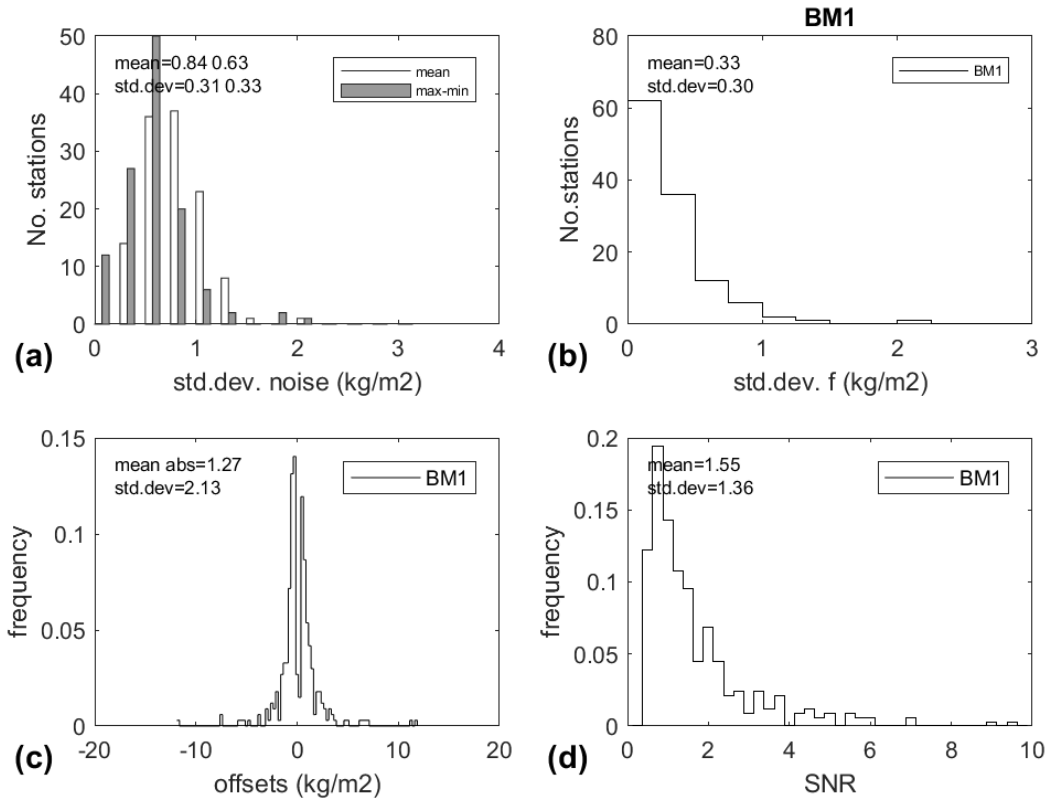


FIGURE 8. Histograms of segmentation results for the final method with selection criterion BM1: (a) Number of stations with respect to the estimated standard deviation of the noise (mean and max-min of the 12 monthly values); (b) Number of stations with respect to the standard deviation of the estimated function; (c) Distribution of offsets of detected change-points; (d) Distribution of SNR of detected change-points..

1999-07-29, type: R) and is validated. Variant (c) gives no detection, the conservative option is selected by BM1 (option (ii) discussed above).

- In the case of DUBO, variants (a) and (c) detect two change-points at almost the same position but not exactly. Both are located close to known changes but only one is validated for variant (a) (detected: 1999-05-07, known: 1999-05-26, type: R). The second one is located 34 days from a known change for variant (a) and 148 days for variant (c). Though variant (c) works not bad, it is not as accurate as variant (a) because the periodic bias is neglected. Variant (d) has 4 detections which actually consist in 2 change-points, each being associated with an outlier. Although the periodic bias is modelled here, both change-points are quite badly located and thus not validated.
- Finally for MCM4 the signal has very marked inhomogeneities in the form of several abrupt changes but also non-stationary oscillations. The abrupt changes are well captured by variant (a) who detects 5 change-points among which 4 are validated (types are in chronological order: R, R, P, P). The non-stationary oscillations are only partly modelled by the periodic function. This is a special case where even the model used in variant (a) is not well adapted to such oscillations. This result advocates for an improvement of the functional basis. In that case, variant (c) works quite well too and leads to almost the same detections as variant (a), but only the two P changes are validated. Variant (d) on the other hand over-estimates the number of change-points to better fit the non-stationary

oscillations but with detections of outliers. The four same change-points are validated as with variant (a) but the fitted means are quite different.

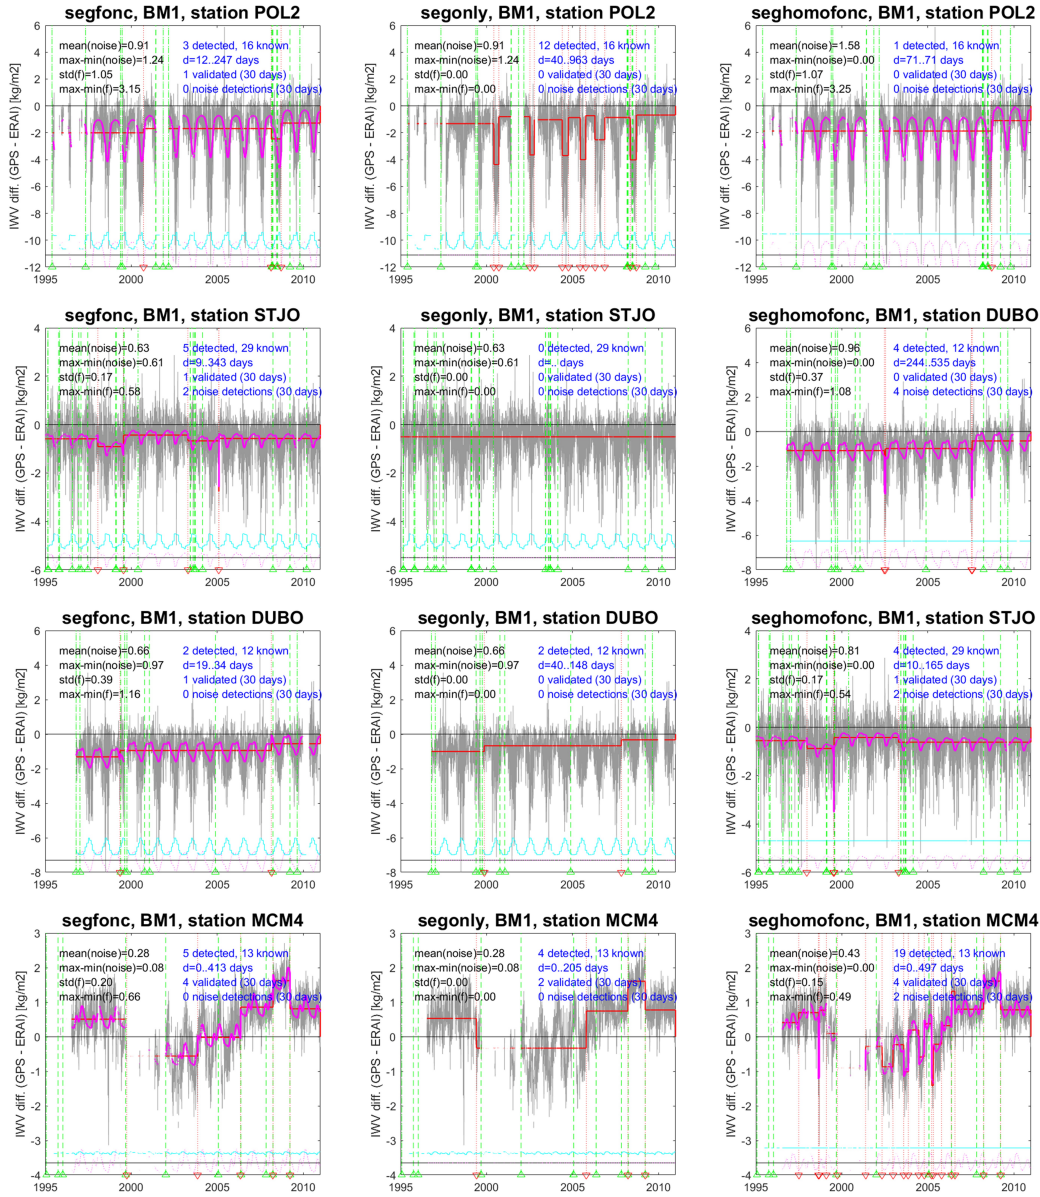


FIGURE 9. Examples of results obtained with variants (a), (c), and (d) from left to right, for four different stations: POL2, STJO, DUBO, and MCM4 (from top to bottom). The content of the plots is similar to Fig. 1(b). The text inserted at the top left of the plots reports the mean standard deviation of the noise, the variation (max-min) of the standard deviation of the noise, the standard deviation of the periodic bias function, and the variation (max-min) of the periodic bias function. The text in blue reports the total number of detections and of known changes, the minimum and maximum distance between detected change-points and the nearest known changes, the number of validated detections, and the number of noise detections.

Among the 70 validated change-points found by BM1 in the case of variant (a) there are 53 R, 16 A, 7 D, and 13 P types (note that these numbers don't sum up to 70 because in many cases the changes involve several types). We find here that receiver changes are the most frequent explanation for inhomogeneities. This is not surprising since they are the most frequent change-type occurring at GNSS stations. However, this is in contrast with Ning *et al.* (2016a)'s results who did not consider receiver changes at all. About 70% of the receiver changes documented in the IGS sitelogs actually refer to firmware updates which don't have much impact on the observations as long as they don't involve a change in the minimum elevation cutoff angle. Hardware changes on the other hand are more prone to have an impact. We performed a quality control based on the observation files with TEQC software (Estey and Meertens, 1999) and found that in many cases hardware changes lead to changes in the multipath diagnostic parameters and in some occasions in the percentage of observations. Receiver changes that have an impact are e.g. found at station STJO on 1999-08-06 (from ROGUE_SNR_8000 to AOA_SNR_12_ACT) and at station MCM4 on 2002-01-03 (from ROGUE_SNR_8000 to AOA_SNR_12_ACT) and on 2006-05-19 (from AOA_SNR_12_ACT to ASHTECH_ZXII3). At MCM4, strong oscillations are found in the multipath diagnostics (mp1 and mp2) during the AOA_SNR_12_ACT period, similar to those seen in the IWV differences (Figure 9). This reveals a malfunctioning of the GNSS equipment also associated with a jump in the mean signal at the beginning and at the end of that period.

4.4. Comparison with Ning *et al.* (2016a). Similar to this study, Ning *et al.* (2016a) analyzed the homogeneity of GNSS-ERAI IWV differences for a global network of 101 GNSS sites with a least 15 years of observations. Their series were used with monthly sampling whereas here we used daily sampling. They used the PMTred test (Wang, 2008) to detect abrupt changes in the mean IWV difference but this model does neither include a periodic bias nor a monthly varying variance. They detected a total of 62 change-points affecting 47 stations among which 45 detections were attributed to the GNSS series, 16 to ERAI, and 1 was undetermined. Their attribution method was based on the comparison of the GNSS candidate series to two or three references series (ERAI, another nearby GNSS series, and/or a nearby VLBI series). Consistency between the two or three detected offsets was used to attribute the change-points to GNSS and disagreement to ERAI (by default). They also validated 13 detections with the GNSS metadata, but they included only antenna, radome, and known microwave absorbing material changes. Their validation window was +/- 6-month wide, i.e. much larger than our +/- 30-day window. We reviewed their validations for 42 of their sites for which we had metadata information from the IGS sitelogs including in our case receiver changes. Using the same 6-month window, we found that 10 out of their 12 undocumented GNSS detections can actually be explained with receiver changes and 2 with receiver+antenna changes (the latter were surprisingly missing in their analysis). Six of these changes agreed with the metadata within 2 months or less. We also found that 5 out of 15 of their change-points attributed to ERAI coincide actually with 2 GNSS receiver changes and 3 antenna changes. Finally, inspection of the GNSS-ERAI IWV difference time series suggests that many of their undocumented detections may be due to outliers and gaps in the time series. This suggests that the implementation of the PMTred test is quite sensitive to fluctuations in the noise, a property similar to that of variant (d) discussed in the previous sub-section.

The comparison of our results for variant (a) with Ning *et al.* (2016a)'s results for 31 common stations which have change-points leads to the following conclusions: (i) our method detects nearly twice more change-points than PMTred (107 vs. 43), (ii) among 32 PMTred detections attributed to GNSS, about 1/3rd coincide with ours within +/- 2 months, 1/3rd within 2-6 months and 1/3rd within more than 6 months, (iii) among 11 PMTred detections attributed to ERAI, 4 change-points coincide with ours within +/- 1 month (the others being about 6 months or more apart) and none of them can actually be explained by GNSS changes (even involving receiver changes). Inspection of the IWV differences and the TEQC diagnostics confirms that the 4 change-points attributed to ERAI cannot be explained by changes in the GNSS time series, i.e. they may truly be due to ERAI; these are: GODE (1998-08-06), HOB2 (2006-06-10), and WUHN (1999-02-14 and 2006-09-27). The latter change-point was already mentioned by Parracho *et al.* (2018) as being due to

a change in radiosonde data from the station at the city of Wuhan, China, being assimilated in ERAI.

5. DISCUSSION AND CONCLUSIONS

In this paper we presented a new segmentation method for the detection of abrupt changes in the mean of geophysical time series including a periodic bias and heterogeneous variance. The results on simulated data showed that the segmentation results (position and amplitude of change-points) are sensitive to the choice of the function basis used to model the periodic bias and to the initialisation of the iterative procedure in which the function and segmentation parameters are estimated. Several model selection criteria were tested. The criterion proposed by Birgé and Massart (2001) and the modified BIC proposed by Zhang and Siegmund (2007) appeared to have good properties. The criterion of Lavielle (2005) appears rather unstable with large dispersion in the number of detected change-points.

When applied to real data (GNSS minus ERAI IWV series), the modified BIC's results were very disappointing (strong over-estimation of the number of change-points), certainly due to the fact that it is derived in the case of a normal distribution and a homoscedastic variance case. In fact, all the considered model selection criteria are based on these assumptions, but according to our experience mBIC is much more sensitive to deviations from the normal distribution.

We tested several variants of the method with the real data and found that accounting for a monthly variance and a period bias improved clearly the detection, although this method has some tendency to detect outliers due to noise spikes (about 20% of the detections). A proper outlier detection method has to be developed, e.g. based on the SNR, in order to reject these detections.

In addition, future improvements of the proposed method would be: (i) to consider other models for the function f since it was found that in some cases like at station MCM4 a simple periodic function is not adequate, (ii) to take the serial correlation in the data into account. The first point can be handled by an estimation of the function f using a non-parametric approach. The second point can be developed by following the approach of Chakar *et al.* (2017) who proposed to model the temporal correlation using an autoregressive process of order 1. These authors also proposed a two-stage whitening inference strategy that allows the use of the DP algorithm and find the exact maximum likelihood solution.

DATA STATEMENT

The GNSS IWV data are available from <https://doi.org/10.14768/06337394-73a9-407c-9997-0e380dac5591>. (last access: April 2020; (Bock , 2016)).

ERA-Interim data are available from <https://www.ecmwf.int/en/forecasts/datasets/archive-datasets/reanalysis-datasets/era-interim> (last access: April 2020; (Dee *et al.* , 2011)).

ACKNOWLEDGEMENTS

This work was developed in the framework of the VEGA project and supported by the CNRS program LEFE/INSU. This paper is IPGP contribution number 4136. The contribution of the third author has been conducted as part of the project Labex MME-DII (ANR11-LBX-0023- 01) and within the FP2M federation (CNRS FR 2036).

REFERENCES

- ARDIA, D., DUFAYS, A. and CRIADO, C. O. (2019). Frequentist and bayesian change-point models: A missing link.
- ARLOT, S. and MASSART, P. (2009). Data-driven calibration of penalties for least-squares regression. *J. Mach. Learn. Res.* **10** 245–279.
- AUGER, I. E. and LAWRENCE, C. E. (Jan, 1989). Algorithms for the optimal identification of segment neighborhoods. *Bulletin of Mathematical Biology.* **51** (1) 39–54.
- BAI, J. and PERRON, P. (2003). Computation and analysis of multiple structural change models. *Journal of Applied Econometrics.* **18** (1) 1–22.

- BERTIN, K., COLLILIEUX, X., LEBARBIER, E. and MEZA, C. (2017). Semi-parametric segmentation of multiple series using a dp-lasso strategy. *Journal of Statistical Computation and Simulation*. **87** (6) 1255–1268.
- BEVIS, M., BUSINGER, S., HERRING, T. A., ROCKEN, C., ANTHES, R. A. and WARE, R. H. (1992). Gps meteorology: Remote sensing of atmospheric water vapor using the global positioning system. *JOURNAL OF GEOPHYSICAL RESEARCH*. **97** (15) 787–801.
- BIRGÉ, L. and MASSART, P. (2001). Gaussian model selection. *Journal of the European Mathematical Society*. **3** 203–268.
- BOCK, O., (2016). Gps data: Daily and monthly reprocessed iwv data from 120 global gps stations, version 1.2.
- BOCK, O., BOSSER, P., BOURCY, T., DAVID, L., GOUTAIL, F., HOAREAU, C., KECKHUT, P., LEGAIN, D., PAZMINO, A., PELON, J., PIPIS, K., POUJOL, G., SARKISSIAN, A., THOM, C., TOURNOIS, G. and TZANOS, D. (2013). Accuracy assessment of water vapour measurements from in situ and remote sensing techniques during the demevap 2011 campaign at ohp. *Atmospheric Measurement Techniques*. **6** (10) 2777–2802.
- BOCK, O., COLLILIEUX, X., GUILLAMON, F., LEBARBIER, E. and PASCAL, C. (2018). A breakpoint detection in the mean model with heterogeneous variance on fixed time-intervals. *Statistics and Computing*. **63** (1) 22–32.
- BOCK, O. and PARRACHO, A. (2019). Consistency and representativeness of integrated water vapour from ground-based gps observations and era-interim reanalysis. *Atmos. Chem. Phys.* **19** 9453–9468.
- CAUSSINUS, H. and MESTRE, O. (2004). Detection and correction of artificial shifts in climate series. *Journal of the Royal Statistical Society: Series C (Applied Statistics)*. **53** (3) 405–425.
- CHAKAR, S., LEBARBIER, E., LÉVY-LEDUC, C., ROBIN, S. *et al.* (2017). A robust approach for estimating change-points in the mean of an AR(1) process. *Bernoulli*. **23** (2) 1408–1447.
- COSTA, A. C. and SOARES, A. (Apr, 2009). Homogenization of climate data: Review and new perspectives using geostatistics. *Mathematical Geosciences*. **41** (3) 291–305.
- DEE, D. P., UPPALA, S., SIMMONS, A., BERRISFORD, P., POLI, P., KOBAYASHI, S., ANDRAE, U., BALMASEDA, M., BALSAMO, G., BAUER, D. P. *et al.* (2011). The era-interim reanalysis: Configuration and performance of the data assimilation system. *Quarterly Journal of the royal meteorological society*. **137** (656) 553–597.
- EASTERLING, D. and PETERSON, T. (1995). A new method for detecting undocumented discontinuities in climatological time series. *Int. J. Climatol.* **15** 369–377.
- ESTEY, L. and MEERTENS, C. (1999). Teqc: the multi-purpose toolkit for gps/glonass data. *GPS Solutions*. **3** 42–49.
- GAZEAUX, J., LEBARBIER, E., COLLILIEUX, X. and MÉTIVIER, L. (2015). Joint segmentation of multiple gps coordinate series. *Journal de la Société Française de Statistique*. **156** (4) 163–179.
- JONES, P. D., RAPER, S. C. B., BRADLEY, R. S., DIAZ, H. F., KELLYO, P. M. and WIGLEY, T. M. L. (1986). Northern hemisphere surface air temperature variations: 1851–1984. *Journal of Climate and Applied Meteorology*. **25** (2) 161–179.
- KILLICK, R., FEARNHEAD, P. and ECKLEY, I. A. (2012). Optimal detection of changepoints with a linear computational cost. *Journal of the American Statistical Association*. **107** (500) 1590–1598.
- LAVIELLE, M. (2005). Using penalized contrasts for the change-point problem. *Signal Processing*. **85** (8) 1501–1510.
- LEBARBIER, E. (2005). Detecting multiple change-points in the mean of Gaussian process by model selection. *Signal Processing*. **85** 717–736.
- LI, S. and LUND, R. (2012). Multiple changepoint detection via genetic algorithms. *Journal of Climate*. **25** (2) 674–686.
- LU, Q., LUND, R. and LEE, T. C. M. (2010). An mdl approach to the climate segmentation problem. *The Annals of Applied Statistics*. **4** (1) 299–319.
- MAIDSTONE, R., HOCKING, T., RIGAILL, G. and FEARNHEAD, P. (2017). On optimal multiple changepoint algorithms for large data. *Stat. Comput.* **27** 519–533.
- MENNE, M. J. and WILLIAMS, C. N. (2005). Detection of undocumented changepoints using

- multiple test statistics and composite reference series. *Journal of Climate*. **18** (20) 4271–4286.
- NING, T., WICKERT, J., DENG, Z., HEISE, S., DICK, G., VEY, S. and SCHÖNE, T. (2016a). Homogenized time series of the atmospheric water vapor content obtained from the gnss reprocessed data. *Journal of Climate*. **29** (7) 2443–2456.
- NING, T., WANG, J., ELGERED, G., DICK, G., WICKERT, J., BRADKE, M., SOMMER, M., QUEREL, R. and SMALE, D. (2016b). The uncertainty of the atmospheric integrated water vapour estimated from gnss observations. *Atmos. Meas. Tech.* **9** (1) 79–92.
- PARRACHO, A. C., BOCK, O. and BASTIN, S. (2018). Global iwv trends and variability in atmospheric reanalyses and gps observations. *Atmospheric Chemistry and Physics*. **18** (22) 16213–16237.
- PETERSON, T. C., EASTERLING, D. R., KARL, T. R., GROISMAN, P., NICHOLLS, N., PLUMMER, N., TOROK, S., AUER, I., BOEHM, R., GULLETT, D. *et al.* (1998). Homogeneity adjustments of in situ atmospheric climate data: a review. *International Journal of Climatology: A Journal of the Royal Meteorological Society*. **18** (13) 1493–1517.
- PICARD, F., ROBIN, S., LAVIELLE, M., VAISSE, C. and DAUDIN, J.-J. (Feb, 2005). A statistical approach for array cgh data analysis. *BMC Bioinformatics*. **6** (1) 27.
- REEVES, J., CHEN, J., WANG, X. L., LUND, R. and LU, Q. Q. (2007). A review and comparison of changepoint detection techniques for climate data. *Journal of Applied Meteorology and Climatology*. **46** (6) 900–915.
- RIGAILL, G. (2015). A pruned dynamic programming algorithm to recover the best segmentations with 1 to k_m change-points. *Journal de la Société Française de Statistique*. **156** (4) 180–205.
- RISSANEN, J. (1978). Modelling by the shortest data description. *Automatica*. **14** 465–471.
- ROUSSEUW, P. J. and CROUX, C. (1993). Alternatives to the median absolute deviation. *Journal of the American Statistical Association*. **88** (424) 1273–1283.
- SZENTIMREY, T. (2008). Development of mash homogenization procedure for daily data. proceedings of the fifth seminar for homogenization and quality control in climatological databases. *WCDMP-No. 71*. 123–130.
- TRUONG, C., OUDRE, L. and VAYATIS, N. (2020). Selective review of offline change point detection methods. *Signal Processing*. **167** 107299.
- VAN MALDEREN, R., BRENOT, H., POTTIAUX, E., BEIRLE, S., HERMANS, C., DE MAZIERE, M., WAGNER, T., DE BACKER, H. and BRUYNINX, C. (08, 2014). A multi-site intercomparison of integrated water vapour observations for climate change analysis. **7**.
- VARADHAN, R. and ROLAND, C. (2008). Simple and globally convergent methods for accelerating the convergence of any em algorithm. *Scandinavian Journal of Statistics*. **35** (2) 335–353.
- VENEMA, V. K. C., MESTRE, O., AGUILAR, E., AUER, I., GUIJARRO, J. A., DOMONKOS, P., VERTACNIK, G., SZENTIMREY, T., STEPANEK, P., ZAHRADNICEK, P., VIARRE, J., MÜLLER-WESTERMEIER, G., LAKATOS, M., WILLIAMS, C. N., MENNE, M. J., LINDAU, R., RASOL, D., RUSTEMEIER, E., KOLOKYTHAS, K., MARINOVA, T., ANDRESEN, L., ACQUAOTTA, F., FRATIANNI, S., CHEVAL, S., KLANCAR, M., BRUNETTI, M., GRUBER, C., PROHOM DURAN, M., LIKSO, T., ESTEBAN, P. and BRANDSMA, T. (2012). Benchmarking homogenization algorithms for monthly data. *Climate of the Past*. **8** (1) 89–115.
- VEY, S., DIETRICH, R., FRITSCHÉ, M., RÜLKE, A., STEIGENBERGER, P. and ROTHACHER, M. (2009). On the homogeneity and interpretation of precipitable water time series derived from global gps observations. **114** (D10).
- WANG, X. L. (2008). Accounting for autocorrelation in detecting mean shifts in climate data series using the penalized maximal t or f test. *Journal of Applied Meteorology and Climatology*. **47** (9) 2423–2444.
- WEATHERHEAD, E. C., REINSEL, G. C., TIAO, G. C., MENG, X., CHOI, D., CHEANG, W., KELLER, T., DELUISI, J., WUEBBLES, D. J., KERR, J. B., MILLER, A. J., OLTMANS, S. J. and FREDERICK, J. E. (1998). Factors affecting the detection of trends: Statistical considerations and applications to environmental data.
- ZHANG, N. R. and SIEGMUND, D. O. (2007). A modified Bayes information criterion with applications to the analysis of comparative genomic hybridization data. *Biometrics*. **63** (1) 22–32.

SUPPLEMENTAL MATERIAL

Summary of the proposed procedure. Figure 10 summarizes the proposed procedure.

Tested alternatives to the proposed procedure. Recall that in our procedure (see Section 2.2), (1) the variances are estimated first; (2) the iterative procedure is initialized by the estimation of f using an unweighted least-square criterion; (3) the function is estimated with a Fourier decomposition of order 4. We tested different variants for these three points:

(1) Updating the variances: we tested a version of the procedure where σ was updated at each iteration of the iterative procedure. The estimated variances are plotted in Figure 11. This option provided slightly more accurate estimates for all the variance (see Figure 3) and the function parameters with very little impact on the segmentation parameters (not shown) compared to our procedure. However, the small changes in variance at each iteration severely slowed down the convergence of the algorithm.

(2) Variants of the initialization: three variants are tested: (a) the segmentation is performed first; (b) f is estimated first using a weighted regression (as in the iterative procedure); (c) f is estimated first using a weighted regression but on the centered signal $y_t - \bar{y}$.

Figure 12 shows the results for option (a). Compared to the results of our procedure (see Figs. 4 and 6), the results are significantly degraded. Especially, the larger d_1 indicates that change-points are badly located. At the beginning, the unmodelled periodic variations present in the signal are captured by the segmentation. The iterative procedure does not change this effect leading naturally to an over-segmentation in addition of the bad estimation of f . This is particularly marked for small values of the noise σ_2 and for the Lav criterion whatever σ_2 .

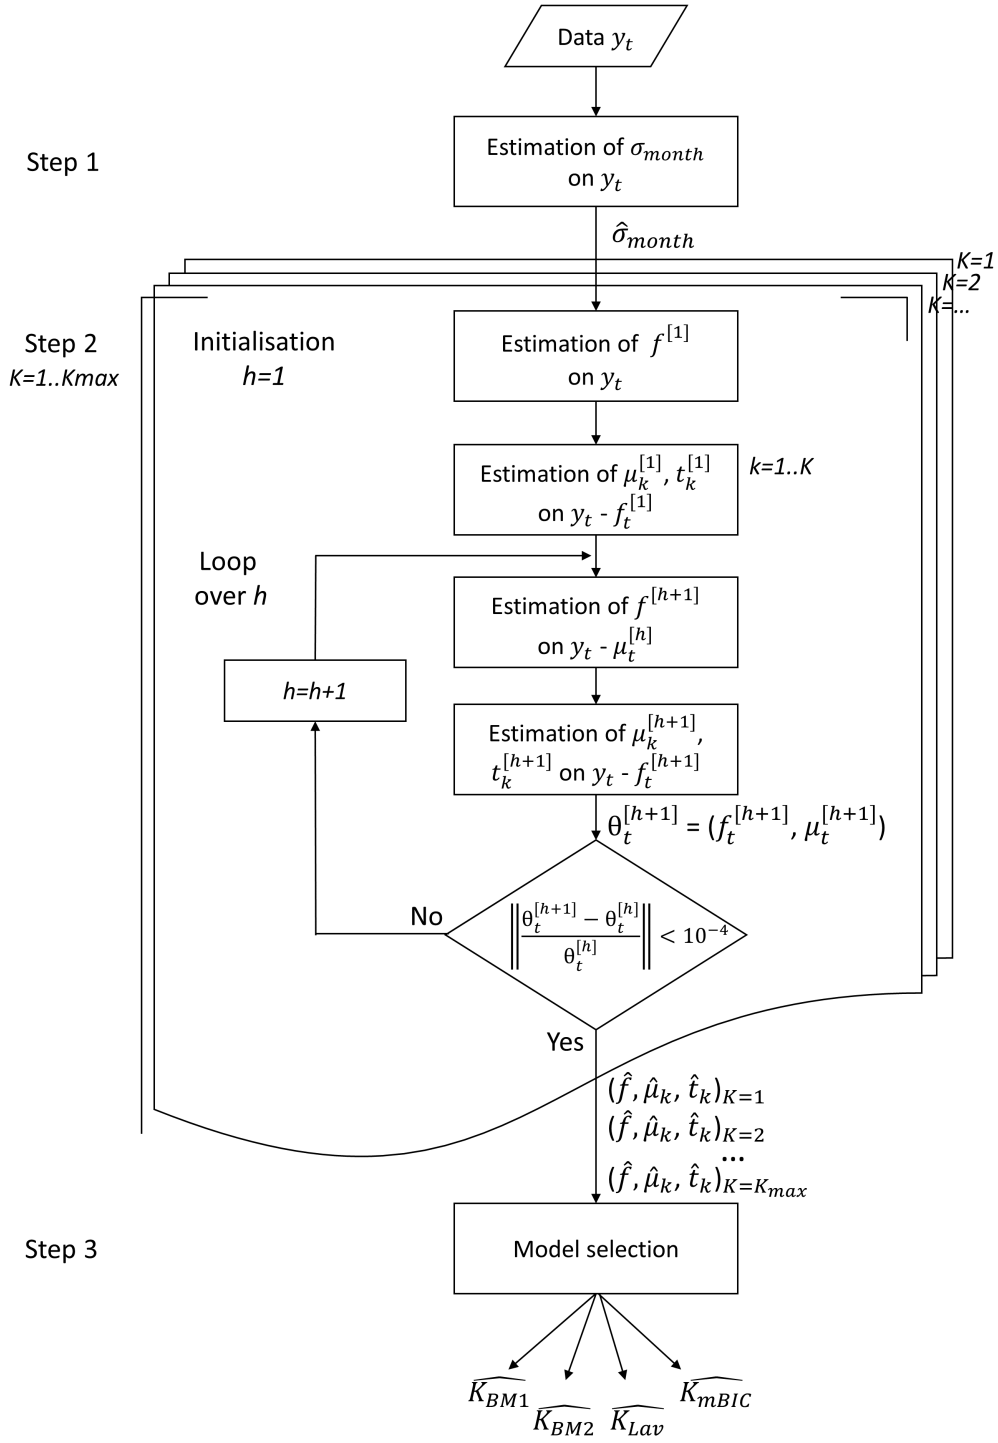
Figure 13 shows the results for option (b). The results are degraded as well but less than previously and mainly for larger σ_2 . This can be explained by the fact that the unmodelled change-points belonging to small variance periods are absorbed by f degrading thus its estimation at this initialization step. And as previously, the iterative procedure does not correct this effect.

The results for option (c) (not shown here) are very similar to those obtained with our initialization procedure. This alternative is equivalent to include a constant term in the linear regression to estimate f . Its estimation is less degraded compared to option (b) and it is correct in the loop.

Our choice of estimating first the function f using an unweighted regression is more flexible in the sense that it does not capture the all segmentation effect at the initialization step allowing thus the iterative procedure to correctly separate the function and the segmentation terms.

(3) Function model: The sensitivity of the procedure to the initialization step discussed above highlights the possible confusion between the function and segmentation. This sensitivity can be further explored by testing different models for f . The idea behind is that simpler models might be less confused with the segmentation making the procedure more accurate in terms of change-point locations. We tested two alternatives: (a) the shape of f is known up to a scaling factor, i.e. $f_t = a_1 \cos(2\pi t/L)$; (b) the statistically significant terms of the Fourier series are selected which have a p-value ≤ 0.001 . Figure 14 and 15 show that the results for these two cases are both consistent and improve the segmentation results compared to our method (see Figure 4 and 6) as expected. Especially, the overall RMSE of the fitted function is strongly reduced. The impact on the positions and amplitudes of the change-points is rather small, however, and the impact in the case of real data is negligible (see Section 4). This test points to the importance of the function model in our method. However, when it comes to real data, the real form of the function is not well known, i.e. the Fourier series of order 4 or even higher may be inadequate. It might thus be useful in a future version of the method to use a more complex base of functions.

FIGURE 10. Schematic of the algorithm.



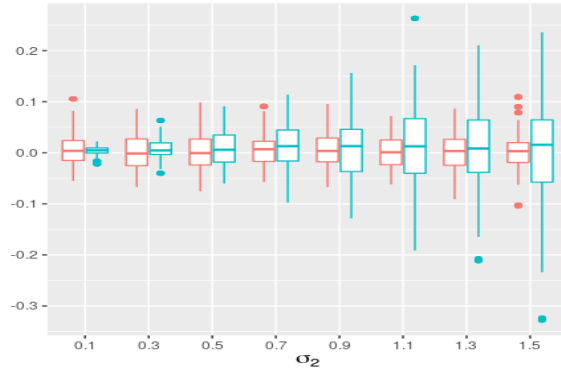


FIGURE 11. Boxplots of standard deviation estimation errors for variant (1): $\hat{\sigma}_1 - \sigma_1^*$ in red and $\hat{\sigma}_2 - \sigma_2^*$ in blue, with $\sigma_1^* = 0.5$ and $\sigma_2^* = 0.1, \dots, 1.5$.

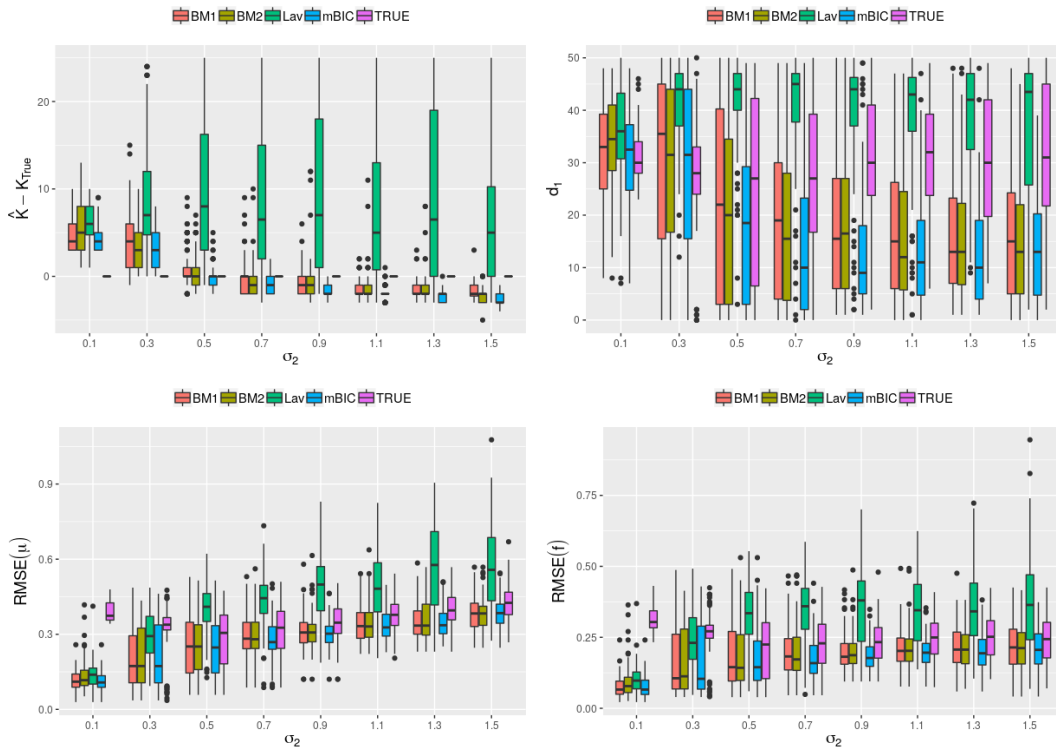


FIGURE 12. Results for variant 2-(a). (a) $\hat{K} - K^*$; (b) first Hausdorff distance d_1 ; (c) $\text{RMSE}(\mu)$; (d) $\text{RMSE}(f)$.

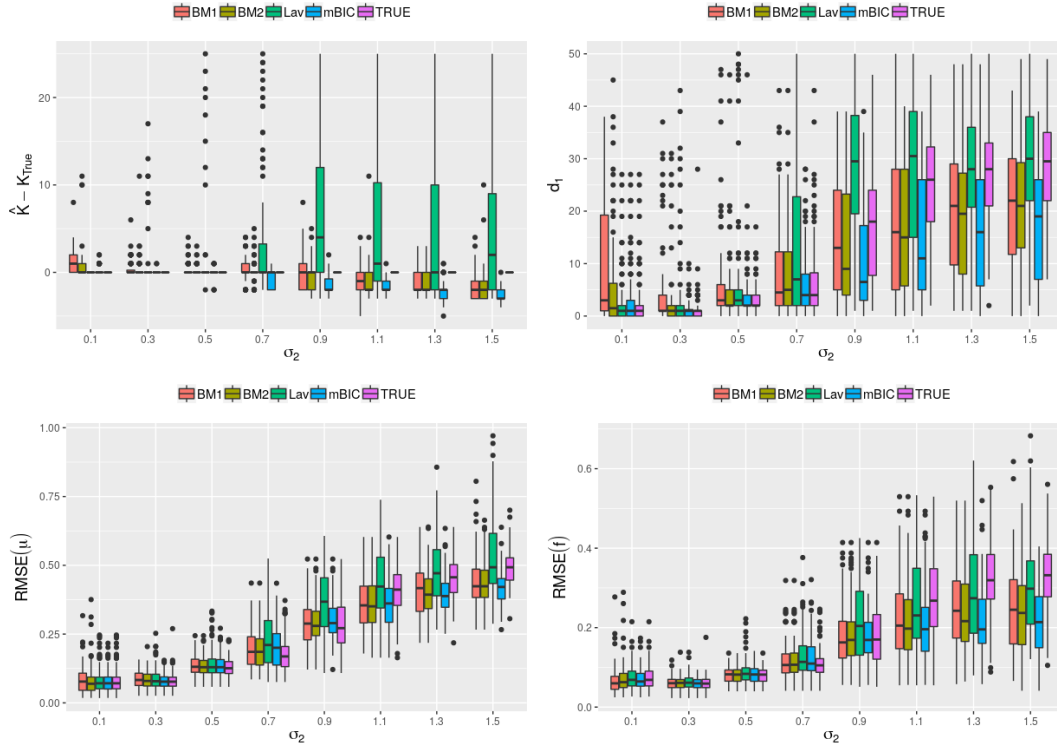


FIGURE 13. Results for variant 2-(b). (a) $\hat{K} - K^*$; (b) first Hausdorff distance d_1 ; (c) $\text{RMSE}(\mu)$; (d) $\text{RMSE}(f)$.

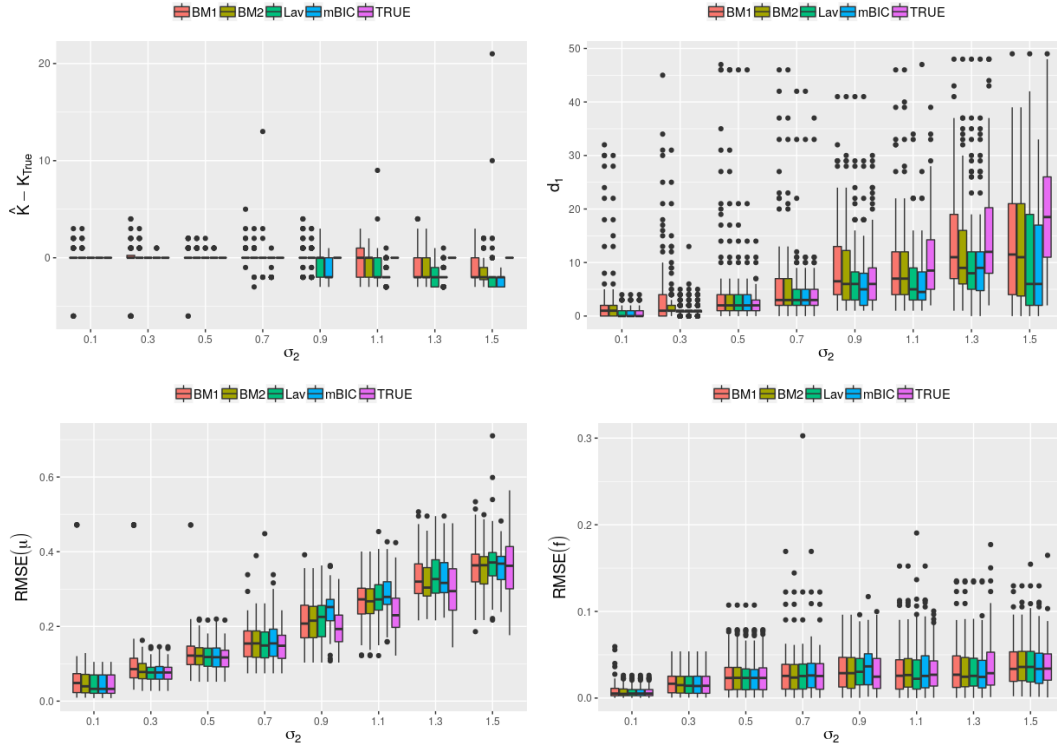


FIGURE 14. Results for variant 3-(a). (a) $\hat{K} - K^*$; (b) first Hausdorff distance d_1 ; (c) $\text{RMSE}(\mu)$; (d) $\text{RMSE}(f)$.

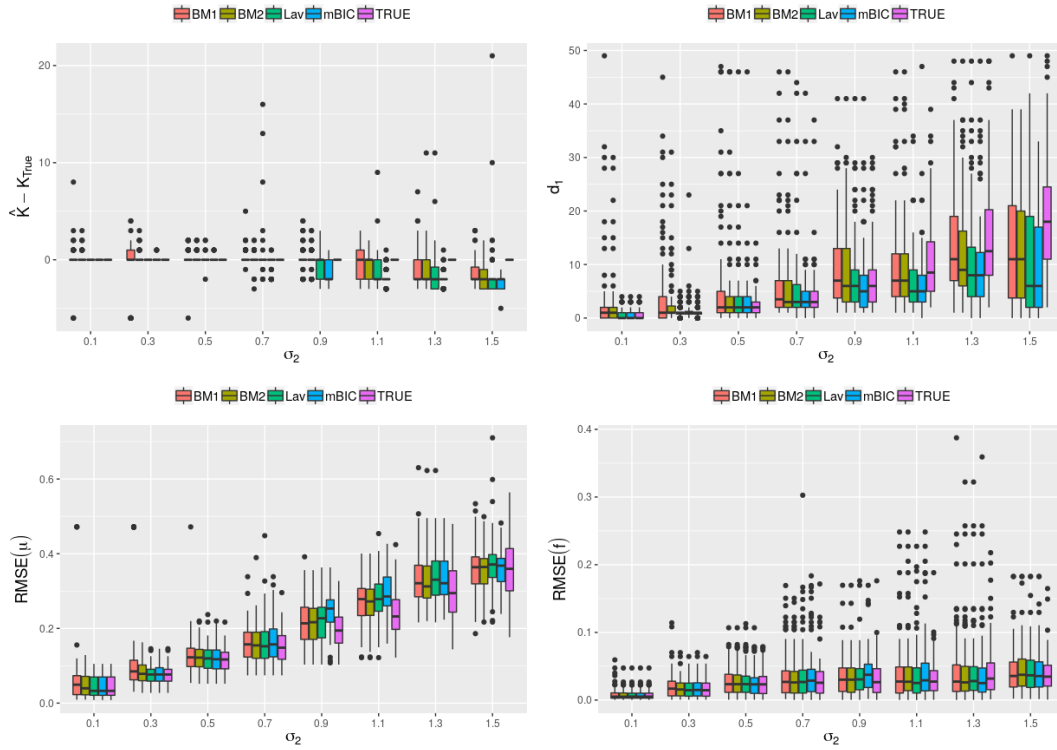


FIGURE 15. Results for variant 3-(b). (a) $\hat{K} - K^*$; (b) first Hausdorff distance d_1 ; (c) $\text{RMSE}(\mu)$; (d) $\text{RMSE}(f)$.