



Seeded intervals and noise level estimation in change point detection: a discussion of Fryzlewicz (2020)

Solt Kovács¹ · Housen Li² · Peter Bühlmann¹

Received: 9 June 2020 / Accepted: 12 June 2020 / Published online: 16 September 2020
© Korean Statistical Society 2020

Abstract

In this discussion, we compare the choice of seeded intervals and that of random intervals for change point segmentation from practical, statistical and computational perspectives. Furthermore, we investigate a novel estimator of the noise level, which improves many existing model selection procedures (including the steepest drop to low levels), particularly for challenging frequent change point scenarios with low signal-to-noise ratios.

Keywords Break points · Fast computation · Model selection · Reproducibility · Seeded binary segmentation · Steepest drop to low levels · Variance estimation · Wild binary segmentation 2

1 Starting remarks

We congratulate Piotr Fryzlewicz for his interesting and stimulating paper entitled “Detecting possibly frequent change-points: Wild Binary Segmentation 2 and steepest-drop model selection” (see Fryzlewicz 2020) and we also thank the Editor of the Journal of the Korean Statistical Society for the invitation to write a discussion! We also take the opportunity to express that we have been very much inspired by Fryzlewicz’ previous work on change point detection. We mention in particular his path-breaking work on wild binary segmentation (Fryzlewicz 2014) which has heavily shaped the thinking on change point detection as well as the narrowest over threshold method (Baranowski et al. 2019). Our remarks and comments should be seen in this light, owing credit to many pioneering ideas from Fryzlewicz. His current paper proposes a new change point detection algorithm (WBS2) that might be thought of as a

This comment refers to the article <https://doi.org/10.1007/s42952-020-00060-x>.

✉ Solt Kovács
solt.kovacs@stat.math.ethz.ch

¹ Seminar for Statistics, ETH Zurich, Zurich, Switzerland

² Institute for Mathematical Stochastics, University of Göttingen, Göttingen, Germany

hybrid between binary segmentation (BS, Vostrikova 1981) and wild binary segmentation (WBS, Fryzlewicz 2014), as well as a novel model selection procedure (steepest drop to low levels, SDLL). While recently several algorithms have been proposed for change point detection, some of which are computationally very efficient, less attention is paid to model selection which in our point of view remains a difficult task. We thus welcome and appreciate new methods such as SDLL that contribute to model selection. In the following we would like to point to a number of alternatives and modifications that might lead to improved performance in terms of stability/reproducibility, estimation error, range of extendability and computational speed.

2 A summary of seeded binary segmentation (SeedBS, Kovács et al. 2020a)

The computational cost for evaluating the CUSUM statistics in M random intervals within WBS is of order $O(MT)$, with the worst case $M = O(T^2)$ (for very short minimal segment length as e.g. in frequent change point scenarios or when aiming to obtain the “complete solution path” defined by Fryzlewicz 2020). In Kovács et al. (2020a), we proposed a deterministic construction of “seeded” intervals that have a total length $O(T \log(T))$ at most and thus allow fast evaluation of the test statistics in all scenarios (as opposed to the worst case $O(T^3)$ for the random intervals in WBS). The idea is roughly to start with the full range of observations in the first layer and sequentially split the intervals from the previous layer into a left, right and an overlapping middle interval according to a chosen “decay” parameter (e.g. $(1, T)$, $(1, T/2)$, $(T/4, 3T/4)$, $(T/2, T)$, $(1, T/4)$, $(T/8, 3T/8)$, etc.) and stopping at the layer where the length of intervals drops to a pre-defined minimal length m . Seeded binary segmentation (SeedBS) with greedy selection resembles WBS, while SeedBS with the narrowest over threshold (NOT) selection is similar to the NOT method of Baranowski et al. (2019), the difference in both cases is essentially to use seeded intervals instead of random ones. Seeded intervals consist of $O(T)$ intervals efficiently covering various scales, in particular guaranteeing good coverage of each single change point which in turn ensures good estimation performance (e.g. minimax optimality for the Gaussian change in mean setup for SeedBS with NOT selection for a suitable threshold). For further details on theory, performance in simulations, explanations on where the computational inefficiency of random intervals comes from and how seeded intervals overcome this, see Kovács et al. (2020a).

3 Adapting SeedBS to frequent change point scenarios

In WBS2 (with $\tilde{M} = 100$) once the current search interval (s, e) has less than 15 observations, all intervals within (s, e) are evaluated, facilitating estimation (and model selection) in frequent change point scenarios when the true segments are extremely short. A possible adaption to SeedBS in order to match this is to include all intervals below a certain length additionally to the regularly generated seeded intervals. In the experiments further below we added all intervals having less than

10 observations. Alternatively, when aiming to improve results in very challenging scenarios, then instead of the default decay parameter of $2^{1/2}$, one could generate more seeded intervals by changing the decay to e.g. $2^{1/8}$, resulting in an increase of computational time by a factor of approximately four. Another option for improving estimation is an adaptive version (i.e. ASeedBS) of SeedBS that we discuss further below.

4 Some advantages of SeedBS: a fast, stable/replicable and easy to generalize alternative

First of all, when the last layer with the smallest considered scale within seeded intervals includes all intervals of length $m = 2$ observations, one obtains a complete solution path. Hence, unlike for WBS, model selection in frequent change point scenarios is still possible with a fixed threshold. In Fig. 2 estimators number 9 and 10 (“SeedBS_THR_1.0_...”) use the threshold $1.0 \cdot \hat{\sigma} \sqrt{2 \log(T)}$ with the JFNL estimator (from Definition 1 further below) and the MAD based estimator as $\hat{\sigma}$, respectively. When paired with the JFNL estimator, this fixed thresholding performs nearly as good as the best competitors (e.g. SDLL) in the lower noise level $\sigma = 0.3$. For $\sigma = 0.45$ the best performing SDLL method performs better in general in estimating the number of change points than the fixed threshold (which is not surprising as the fixed threshold is not adapted to the data). However, as an advantage, the fixed threshold has smaller variance than the SDLL based model selection. Second, SeedBS can be combined with many model selection approaches, in particular also with the SDLL model selection. The advantage here is that the deterministic scheme utilized in the seeded intervals leads to reproducible results unlike WBS2+SDLL, where the estimated number of change points for various runs on the same data can show high variability/instability for high noise levels due to the randomness of the intervals. Aggregation over multiple WBS2 runs and taking the “median run” as proposed by Fryzlewicz (2020) might help to gain more stability, but does not fully eliminate the issue while increasing computational efforts. Third, a SeedBS-BS hybrid can be considered (similar to WBS2 being a hybrid between BS and WBS) by generating new seeded intervals specifically in between previously found change points and thus adapted to the information on the signal obtained in earlier iterations of the procedure. This adaptive seeded binary segmentation (ASeedBS) might perform slightly better than SeedBS in challenging scenarios (at the price of somewhat increased computational times as a new set of seeded intervals are evaluated each time a change point candidate is added). One advantage of ASeedBS compared to WBS2 is reproducibility due to the lack of randomness. Fourth, we would like to mention applicability in models beyond the Gaussian change in mean setup. WBS2 has a greedy nature when selecting the next best split point. This might be limiting when trying to apply WBS2 e.g. in piecewise linear models, where the best single split point over an interval containing multiple change points may not correspond to any of the true change points (see Baranowski et al. 2019 and their proposal of the NOT selection instead of a greedy one to solve this issue). In SeedBS and also ASeedBS greedy selection is not essential and they could be easily combined with

other selection rules, e.g. the NOT selection. Note that evaluating the CUSUM statistic for $\tilde{M} = O(\log T)$ random intervals has the same cost $O(T \log T)$ as evaluating all the $O(T)$ seeded intervals. However, if the spacings between change points are short, it is likely that none of the $\tilde{M} = O(\log T)$ random intervals contain only a single change point as opposed to seeded intervals, for which it is likely that some of the $O(T)$ intervals will be suitable (i.e. contain a single change point). As an illustration of wide applicability, in Kovács et al. (2020a) we used SeedBS to detect changes in (high-dimensional) Gaussian graphical models using a graphical lasso based estimator as in Lonschien et al. (2019). Last, but not least, as a method of $O(T \log(T))$ worst case computational complexity, SeedBS is fast, in particular in our recent Fortran based implementation (planned to be made publicly available) that needed (using the default options) around 0.001 seconds for the `extreme.teeth` signal of Fryzlewicz (2020) and thus about one to two orders of magnitude less than the available R implementation of WBS2. We currently work on the so-called optimistic search technique (Kovács et al. 2020b) reducing the computational times for various methods. Applying this search for SeedBS results in the optimistic seeded binary segmentation (OSeedBS) method, which further reduces the computational complexity to $O(T)$ in the Gaussian change in mean setup when using the NOT selection. From a practical perspective SeedBS and OSeedBS naturally offer the possibility to parallelize the evaluation of the test statistic in various seeded intervals, which can be an advantage when applying it in more expensive high-dimensional setups such as the previously mentioned setup of Lonschien et al. (2019). We tried SeedBS and ASeedBS together with the SDLL model selection on the `extreme.teeth` signal of Fryzlewicz (2020) and the estimation performance for the number of change points was similar to WBS2+SDLL. Much bigger differences seemed to result from a different choice of noise level estimator and thus we present some simulation results for that in the next paragraphs. Overall, in our point of view the main possible advantages of SeedBS compared to WBS and WBS2 are the reproducibility, easy extendability beyond the Gaussian change in mean case through e.g. the NOT selection and, to some extent, the easier derivation of theoretical results.

5 An alternative estimator for the error variance

A commonly used approach for the estimation of noise levels in the Gaussian change in mean setup (with constant variance) is to use a robust scale estimator (e.g. median absolute deviation, MAD) of consecutive differences $(X_2 - X_1, \dots, X_T - X_{T-1})/\sqrt{2}$ so that jumps in the mean of the signal can be treated as outliers. However, in frequent change point scenarios, there can be many of these outliers, potentially leading to severe upwards bias. While Fryzlewicz (2020) wrote that accurate estimation of the noise variance “can be difficult to achieve in frequent change point models”, we would like to show that the task is not impossible if a different strategy is considered. Rather than treating jumps as outliers, we actively consider them in the following sense: when looking at consecutive differences (with lag one), each jump in the signal is counted once, however, when looking at lag two differences,

each jump is counted twice. For example, assuming there is a jump at the k -th observation, out of the four pairs $(X_{k-1} - X_{k-2}, X_k - X_{k-1}, X_{k+1} - X_k, X_{k+2} - X_{k+1})$ only $X_k - X_{k-1}$ includes the jump, but out of four lag two differences $(X_{k-1} - X_{k-3}, X_k - X_{k-2}, X_{k+1} - X_{k-1}, X_{k+2} - X_k)$ two of them, namely $X_k - X_{k-2}$ and $X_{k+1} - X_{k-1}$ include the jump (while the noise is of course contained additionally in both series). The assumption for this “counting” to work is to have minimal segment length of least 2 observations. Overall, the difference in the variance of the lag one and lag two series can be utilized to filter out the jumps as follows.

Definition 1 The jump filtered noise level (JFNL) estimator for the variance σ^2 is defined as

$$\hat{\sigma}_{JFNL}^2 = \max \left\{ 0, 2 \cdot v \left(\frac{X_2 - X_1}{\sqrt{2}}, \dots, \frac{X_T - X_{T-1}}{\sqrt{2}} \right) - v \left(\frac{X_3 - X_1}{\sqrt{2}}, \dots, \frac{X_T - X_{T-2}}{\sqrt{2}} \right) \right\},$$

where $v(Y_1, \dots, Y_l) = \sum_{i=1}^l (Y_i - \bar{Y})^2 / l$ is the empirical variance of Y_1, \dots, Y_l .

Although difference based variance estimation methods have a long history in nonparametric regression (see e.g. Hall et al. 1990), we are not aware of JFNL having been proposed elsewhere. The goal when using the empirical variance is to actively count each of the jumps (which is the core idea when comparing the two lags), whereas robust estimators would treat a part of the jumps as outliers, but possibly to a different extent in the lag one versus the lag two differences such that their comparison becomes difficult. If the true noise level σ^2 is very low compared to the jumps, the difference of the variance terms above can become negative. In such cases the maximum in the definition above kicks in and we threshold at zero in order to have a non-negative variance estimator. While well-defined, estimating zero as a noise level is not particularly useful when applying for model selection purposes. Note however that this typically occurs only in scenarios where change point detection and model selection is an easy task with many available methods, namely in scenarios with low noise levels compared to the sizes of the jumps. In general, lag k differences count jumps k times and thus, potentially even better estimators could be constructed similarly if it is reasonable to assume that each segment contains at least e.g. $k = 3$ or $k = 4$ observations. One could come up with other estimators as well and also try aggregating various estimators for the variance (e.g. via the median) to have an ensemble type estimator which may even outperform individual ones. A related, but slightly different avenue is to try to make the estimator robust. In Fig. 1 we compare JFNL to the MAD estimator (of lag one differences) for various commonly used test signals. In the extreme.teeth signal $\hat{\sigma}_{JFNL}^2$ performs much better than the MAD estimator and also on the other test signals it seems to perform somewhat better. When looking at scenarios with lower noise levels, a higher variance of the JFNL estimator compared to MAD would be visible such that some other estimators could be better suited than JFNL. However, in the shown scenarios

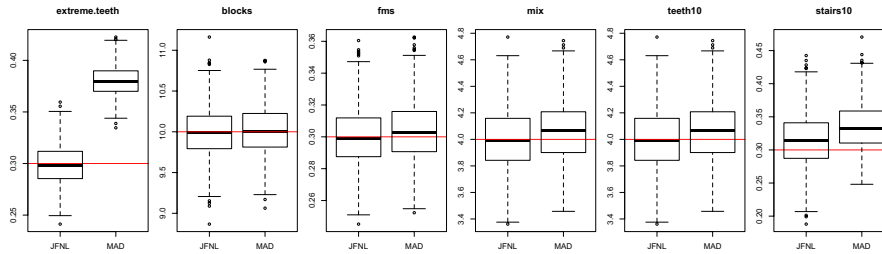


Fig. 1 Boxplots of the JFNL and MAD estimators of the noise level σ in the `extreme.teeth` example of Fryzlewicz (2020) on the very left panel and five other examples from Fryzlewicz (2014) based on 1000 simulations each. The true values of σ are indicated by the red horizontal lines respectively

with comparably high noise (which are the difficult scenarios from a model selection perspective), JFNL seems to offer some advantages. The performance of JFNL for the “stairs10” signal can be further improved by taking $\tilde{v}(Y_1, \dots, Y_l) = \sum_{i=1}^l Y_i^2 / l$ instead of v in Definition 1. As mentioned before, it could be promising to combine strengths of an ensemble of various estimators. There is definitely room for further improvements. It remains to be tested how well $\hat{\sigma}_{JFNL}^2$ performs in examples beyond the ones shown here as well as in real data where the independence and minimal segment length assumptions may not hold.

6 Utilizing the jump filtered noise level estimator $\hat{\sigma}_{JFNL}^2$

For methods/model selection procedures, that rely on an estimate of the noise level one can hope to obtain better results when using a better estimate of the noise level. We illustrate this for some approaches for which changing the default noise level estimate is easy to do within the corresponding R package/implementation. In Fig. 2 we consider the `extreme.teeth` signal with noise level $\sigma = 0.3$ (top) and $\sigma = 0.45$ (bottom). As expected, the JFNL based estimators in general result in more estimated change points than the MAD variants as JFNL does not overestimate the variance. Interestingly, for the lower noise level, FDRSeg (Li et al. 2016) and SeedBS (Kovács et al. 2020a) with a fixed threshold (with JFNL) perform similar to SDLL (with WBS2 or even in combination with SeedBS that we do not show to save space). For higher noise levels, all methods shown in red (JFNL) and black (MAD) tend to find fewer change points. The exception is SDLL with JFNL used for the threshold to determine the “low level” of the drop, which has median estimated number of change points still close to the true number. Hence, a bit surprisingly, it seems like JFNL also improves SDLL methodology in noisy scenarios. SeedBS with a fixed threshold finds fewer change points, but with a smaller variance than SDLL. As the noise level is quite high in this scenario, the question is also what one is aiming for: a good estimate of the number of change points, potentially at the price of including more false positives, or selecting fewer change points and thus also fewer false positives? Note that we did not change the calibration/constants within SDLL when using the JFNL estimator instead of MAD, which might slightly influence the

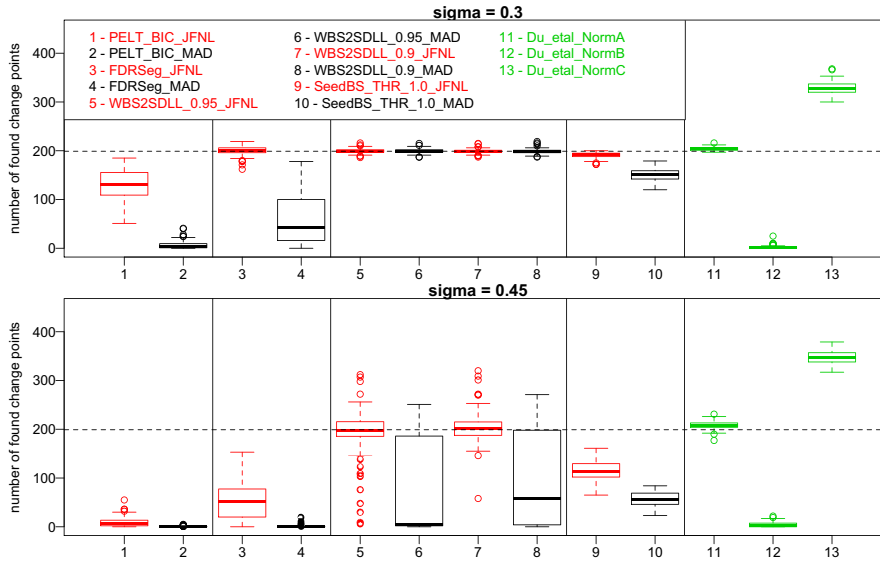


Fig. 2 Boxplots of the number of found change points (based on 100 simulations each) in the `extreme.teeth` example of Fryzlewicz (2020) with noise level $\sigma = 0.3$ (top) and $\sigma = 0.45$ (bottom) for various estimation approaches. Model selection approaches based on the JFNL estimator are in red, MAD based counterparts in black, while green refers to approaches from Du et al. (2016). The true number of change points is indicated by the dashed horizontal line at 199 and the vertical lines group the methods as a visual aid

results. Also, besides the estimated number of change points, other performance measures might be helpful and necessary to understand the detailed behavior of individual methods. Lastly, the green parts of Fig. 2 show estimators from Du et al. (2016) implemented in the `StepSignalMargiLike` package (available from Samuel Kou’s homepage) with the therein proposed three priors (called NormA, NormB and NormC). With the implicit penalization, i.e. default options without us giving any estimate for the noise level, the NormA approach gave very good estimates for the number of change points (with remarkably low variance) even for $\sigma = 0.45$ and thus we considered it worthwhile mentioning. The NormB estimator returned way too few and the NormC way too many candidate change points. In practice, it is a question which prior to choose, moreover, the implementation does not scale well to long time series (around 8 seconds for a time series of length 10^5 and 80 seconds for a time series that is three times longer).

7 BIC based model selection with σ unknown or known?

We note that somewhat confusingly, BIC type criteria are used differently in different publications and software for change point detection. Yao (1988) and Fryzlewicz (2014), for example, minimize $T/2 \cdot \log(\text{MSE}) + \text{penalty}$ (assumption of unknown variance), while dynamic programming approaches (e.g. PELT from Killick et al.

2012) typically minimize $T \cdot \text{MSE} + \text{penalty}$, with MSE denoting the mean squared errors. In the latter approach the penalty term should include the noise level which in this case is assumed to be known. Yet, in practice often an estimate $\hat{\sigma}$ is used in the absence of exact knowledge on σ for which the JFNL estimator could be useful. However, such an approach ignores the uncertainty of the estimated noise level, an issue which is mentioned by Fearnhead and Rigaiil (2020) among many other interesting considerations for practice.

8 Conclusions

We pointed out similarities between SeedBS and WBS2, both being computationally fast methods. The deterministic construction of intervals of SeedBS helps to make results reproducible compared to the randomness inherent in WBS2. Both approaches are generic and thus suitable for many change point detection problems. Our perhaps biased and slight preference for SeedBS also includes that it seems easier to combine seeded intervals with other model selection procedures, such as the NOT method, that are essential in certain change point detection problems beyond the Gaussian change in mean case. We also introduced the JFNL estimator for the noise level and showed in simulations that using JFNL could improve SDLL based model selection in rather noisy scenarios, as well as more traditional model selection approaches relying on a direct estimate of the noise level.

Acknowledgements Solt Kovács and Peter Bühlmann have received funding from the European Research Council (ERC) under the European Union's Horizon 2020 research and innovation programme (Grant agreement No. 786461 CausalStats - ERC-2017-ADG). Housen Li gratefully acknowledges the support of the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) under Germany's Excellence Strategy - EXC 2067/1-390729940.

References

- 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*, 81(3), 649–672.
- Du, C., Kao, C.-L. M., & Kou, S. C. (2016). Stepwise signal extraction via marginal likelihood. *Journal of the American Statistical Association*, 111(513), 314–330.
- Fearnhead, P., & Rigaiil, G. (2020). Relating and comparing methods for detecting changes in mean. *Stat.* <https://doi.org/10.1002/sta4.291>.
- Fryzlewicz, P. (2020). Detecting possibly frequent change-points: Wild Binary Segmentation 2 and steepest-drop model selection. *Journal of the Korean Statistical Society*. <https://doi.org/10.1007/s42952-020-00060-x>.
- Fryzlewicz, P. (2014). Wild binary segmentation for multiple change-point detection. *The Annals of Statistics*, 42(6), 2243–2281.
- Hall, P., Kay, J. W., & Titterton, D. M. (1990). Asymptotically optimal difference-based estimation of variance in nonparametric regression. *Biometrika*, 77(3), 521–528.
- 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.

- Kovács, S., Li, H., Bühlmann, P., & Munk, A. (2020). Seeded binary segmentation: a general methodology for fast and optimal change point detection.
- Kovács, S., Li, H., Haubner, L., Bühlmann, P., & Munk, A. (2020). Optimistic search strategies: change point detection without full grid search. *Working Paper*.
- Li, H., Munk, A., & Sieling, H. (2016). FDR-control in multiscale change-point segmentation. *Electronic Journal of Statistics*, 10(1), 918–959.
- Lonschien, M., Kovács, S., & Bühlmann, P. (2019). Change point detection for graphical models in the presence of missing values. *Journal of Computational and Graphical Statistics*. [arXiv:1907.05409](https://arxiv.org/abs/1907.05409).
- Vostrikova, L. Y. (1981). Detecting disorder in multidimensional random processes. *Soviet Mathematics Doklady*, 24, 55–59.
- Yao, Y.-C. (1988). Estimating the number of change-points via Schwarz' criterion. *Statistics and Probability Letters*, 6(3), 181–189.

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