– Supplementary material –
BayesProject: Fast computation of a projection direction for multivariate changepoint detection

Georg Hahn*, Paul Fearnhead, Idris A. Eckley

Department of Mathematics and Statistics, Lancaster University,
Lancaster, LA1 4YF, U.K.

Contents

The supplementary material provides a variety of further simulations to highlight the performance of BayesProject and its dependence on tuning parameters. In particular, the following sections investigate:

1. The dependence of BayesProject on the length of the series n, on the number of series p in the scenario $p > n$, the proportion of series d/p with a change, the size of the change s, and the changepoint location (Section 1).
2. The dependence on the parameter γ controlling the grid spacing of BayesProject, as well as the robustness of BayesProject to the ratio σ^2/ω^2 used in the computation of the projection direction in eq. (5) in the main article (Section 2).
3. The dependence on the tuning parameter K (Section 3).
4. The performance of BayesProject in the univariate case (Section 4).
5. The proportion of detected changes in BayesProject in the null case when no changepoint exists (Section 5).
6. The asymptotic runtime of Inspect when using a partial SVD (Lanczos iterations) to compute the leading singular vector as opposed to the standard SVD approach (Section 6).
7. An assessment of the accuracy of selected methods with the help of density plots for all scenarios displayed in Tables 1, 2 and 3 in the main article (Section 7).

1 Further simulation studies

This section presents further simulation results comparing the four algorithms investigated in Section 4 of the main article (BayesProject, Inspect, sum-cusum, and max-cusum) with respect

*Corresponding author: ghahn@cantab.net

to their dependence on the series length n, on the number of series p in the scenario $p > n$, the proportion d/p of series with a change, the size s of the change, and the changepoint location. All results in this section are based on 1000 repetitions, and we define sparse ($d/p = 0.05$), moderate ($d/p = 0.2$), and dense ($d/p = 0.5$) scenarios as in the main article. The parameters n, p and s of the changepoint setting are given individually for each figure.

Dependence on n. Figure 1 shows the dependence of the four methods on the length n of the series. Figure 1 shows that Inspect, BayesProject and sum-cusum perform similarly in terms of proportion of estimated changes and average distance to the true changepoint location. Max-cusum is considerably worse. Regarding the accuracy of the estimated projection direction, BayesProject performs best in moderate and dense scenarios, but loses accuracy in the sparse scenario. BayesProject, sum-cusum and max-cusum become more accurate in estimating the projection direction as the changepoint scenario becomes denser. The empirical runtime analysis of all methods indicates a roughly linear dependence on n.

Dependence on p in the scenario $p > n$. Figure 2 repeats the experiments on the dependence on p from the main article for the case $p > n$ while $n = 100$. We see that the estimation of the changepoint (measured via proportion of estimated changes or distance to the true changepoint) is harder for all methods but qualitatively similar to the results reported in the main article. Looking at the narrow range of y-values, we see that the L_2 norm to the ideal changepoint direction is relatively stable for large p in all methods, which is again in line with the results of the main article (BayesProject and sum-cusum estimate the projection direction more accurately than Inspect). Importantly, the difference in asymptotic runtime between BayesProject and Inspect is clearly visible, with Inspect showing a nearly quadratic asymptotic runtime in p while the runtime of BayesProject is linear. Max-cusum performed considerably worse than the other three methods for $p > n$ and in all four cases (proportion, distance, accuracy of estimated projection direction, runtime) its measurements fell outside the range of the ones of the other methods and are thus not visible in Figure 2.

Dependence on d/p. Figure 3 shows the dependence of the four methods on the proportion of series d/p with a change (where $d < p$ is the number of series with a change). The plots display the proportion of estimated changepoints, the average distance to the true changepoint location, the L_2 difference between the ideal projection direction and the estimated projection direction, and the runtime dependence. Figure 3 shows that BayesProject and sum-cusum perform similarly in terms of proportion of estimated changes and average distance to the true changepoint location. Inspect is slightly worse, and max-cusum seems considerably worse. Regarding the accuracy of the estimated projection direction, sum-cusum, BayesProject and max-cusum perform best. Inspect is less accurate. The empirical runtime analysis of all methods shows no dependence on d/p.

Dependence on s. Figure 4 shows the dependence of the four methods on the size s of the change. Figure 4 shows that BayesProject, Inspect and sum-cusum detect an almost equal proportion of estimated changes. BayesProject and sum-cusum perform slightly better than

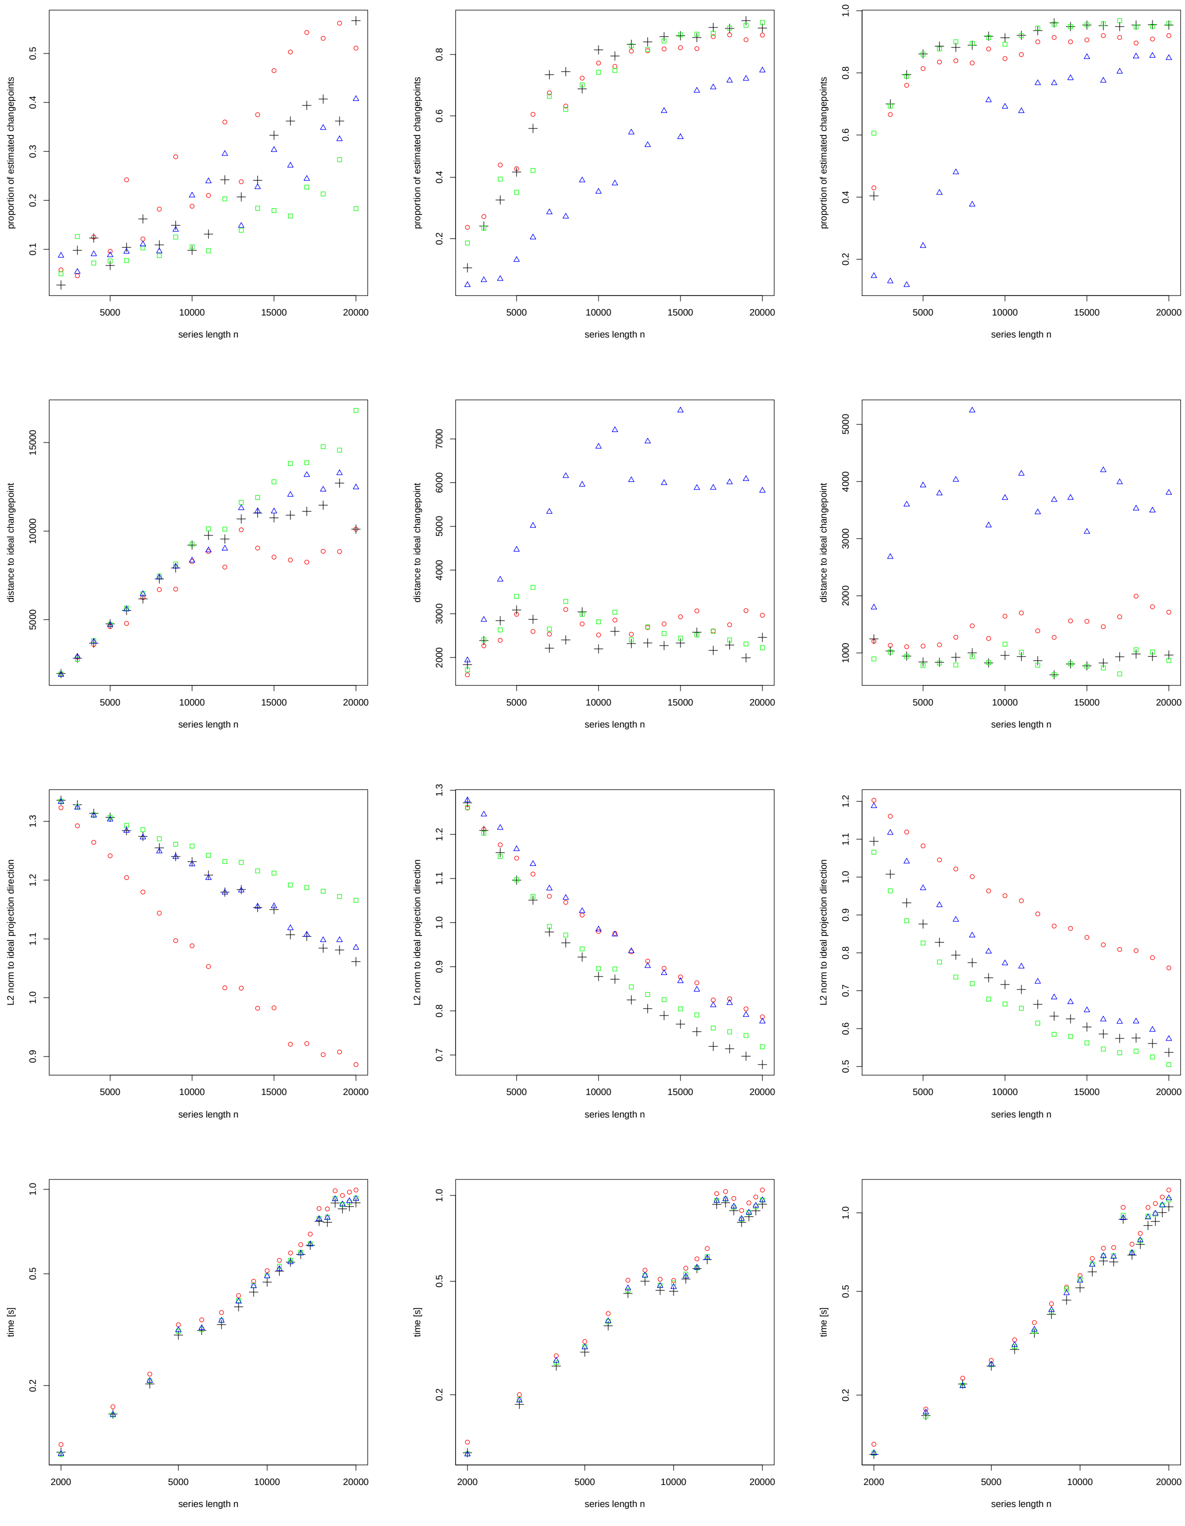

Figure 1: Dependence of BayesProject (black crosses), Inspect (red circles), sum-cusum (green squares), and max-cusum (blue triangles) on the series length n while $p = 100$, $s = 0.05$. Rows show proportion of estimated changes (first row), average distance to the true changepoint location (second row), L_2 distance between the estimated and ideal projection directions (third row), and runtime (fourth row). Columns show sparse (left), moderate (middle), and dense (right) scenarios. Log-log plots in the fourth row. Slope estimates for the moderate scenario (middle column) are 0.88 (Inspect), 0.87 (BayesProject), 0.89 (sum-cusum) and 0.89 (max-cusum).

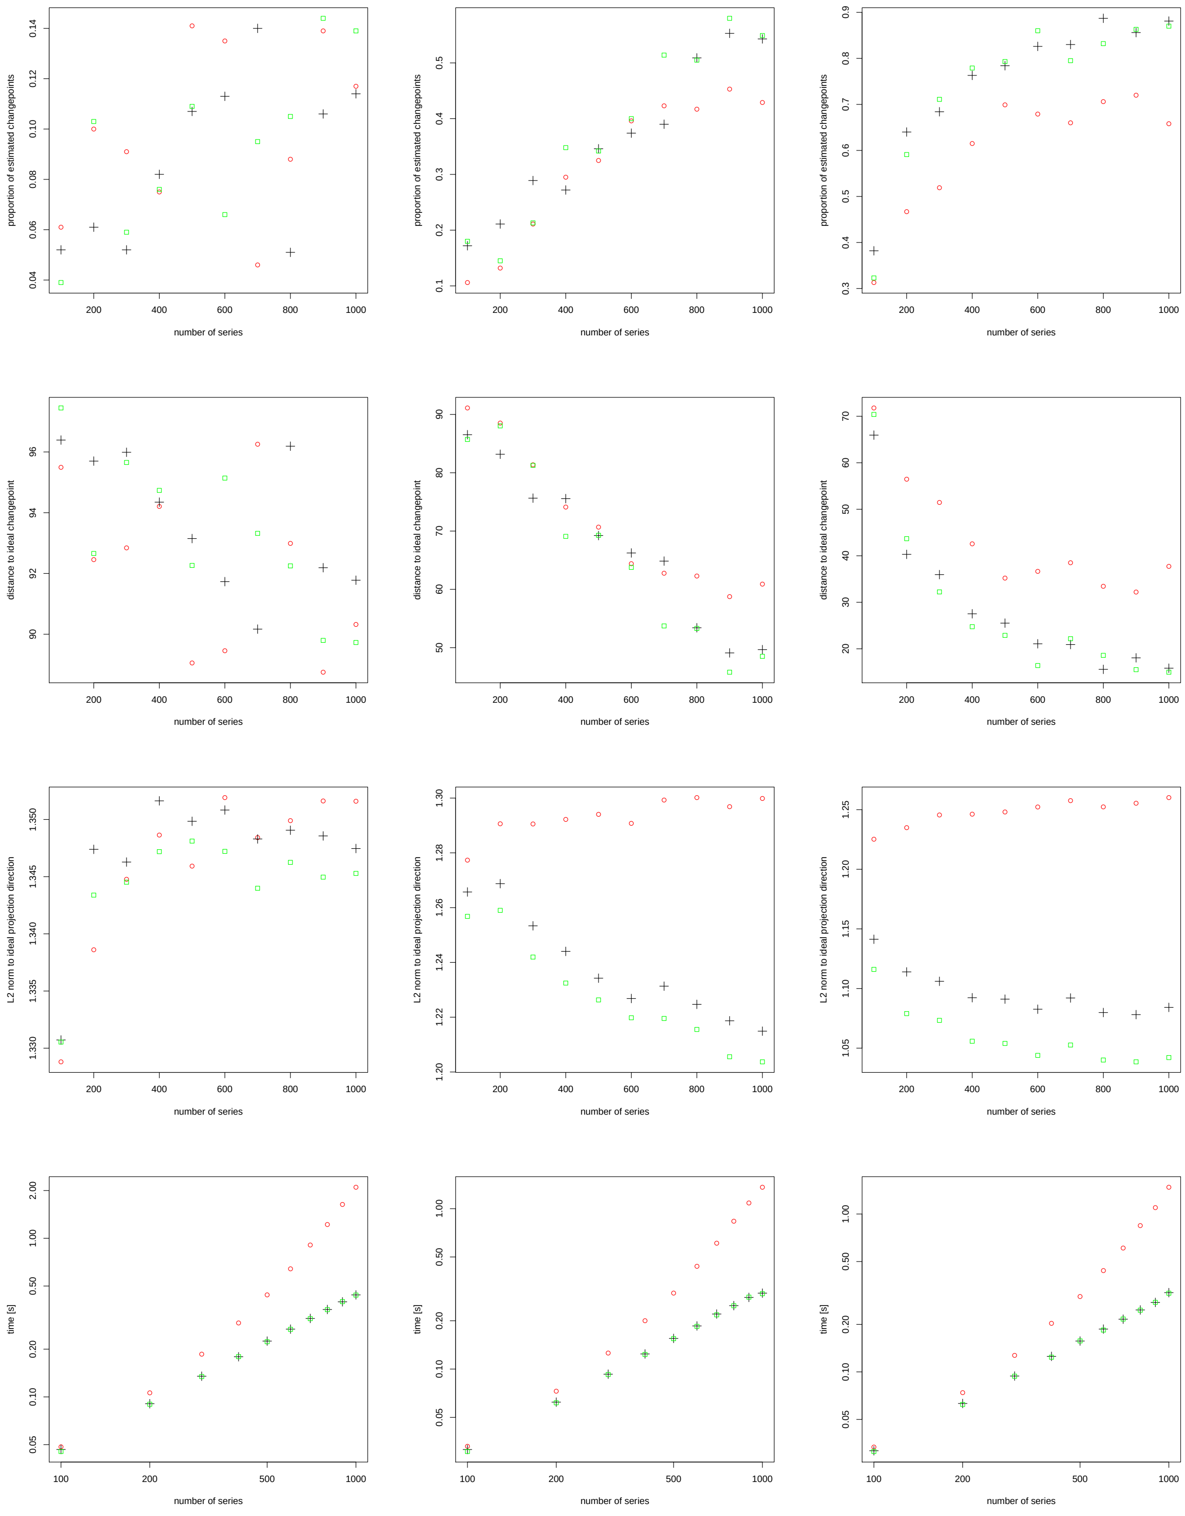

Figure 2: Dependence of BayesProject (black crosses), Inspect (red circles) and sum-cusum (green squares) on the number of series $p > n$ while $n = 100$, $s = 0.2$. Rows show proportion of estimated changes (first row), average distance to the true changepoint location (second row), L_2 distance between the estimated and ideal projection directions (third row), and runtime (fourth row). Columns show sparse (left), moderate (middle), and dense (right) scenarios. Log-log plots in the fourth row. Slope estimates for the moderate scenario (middle column) are 1.62 (Inspect), 0.97 (BayesProject) and 0.98 (sum-cusum).

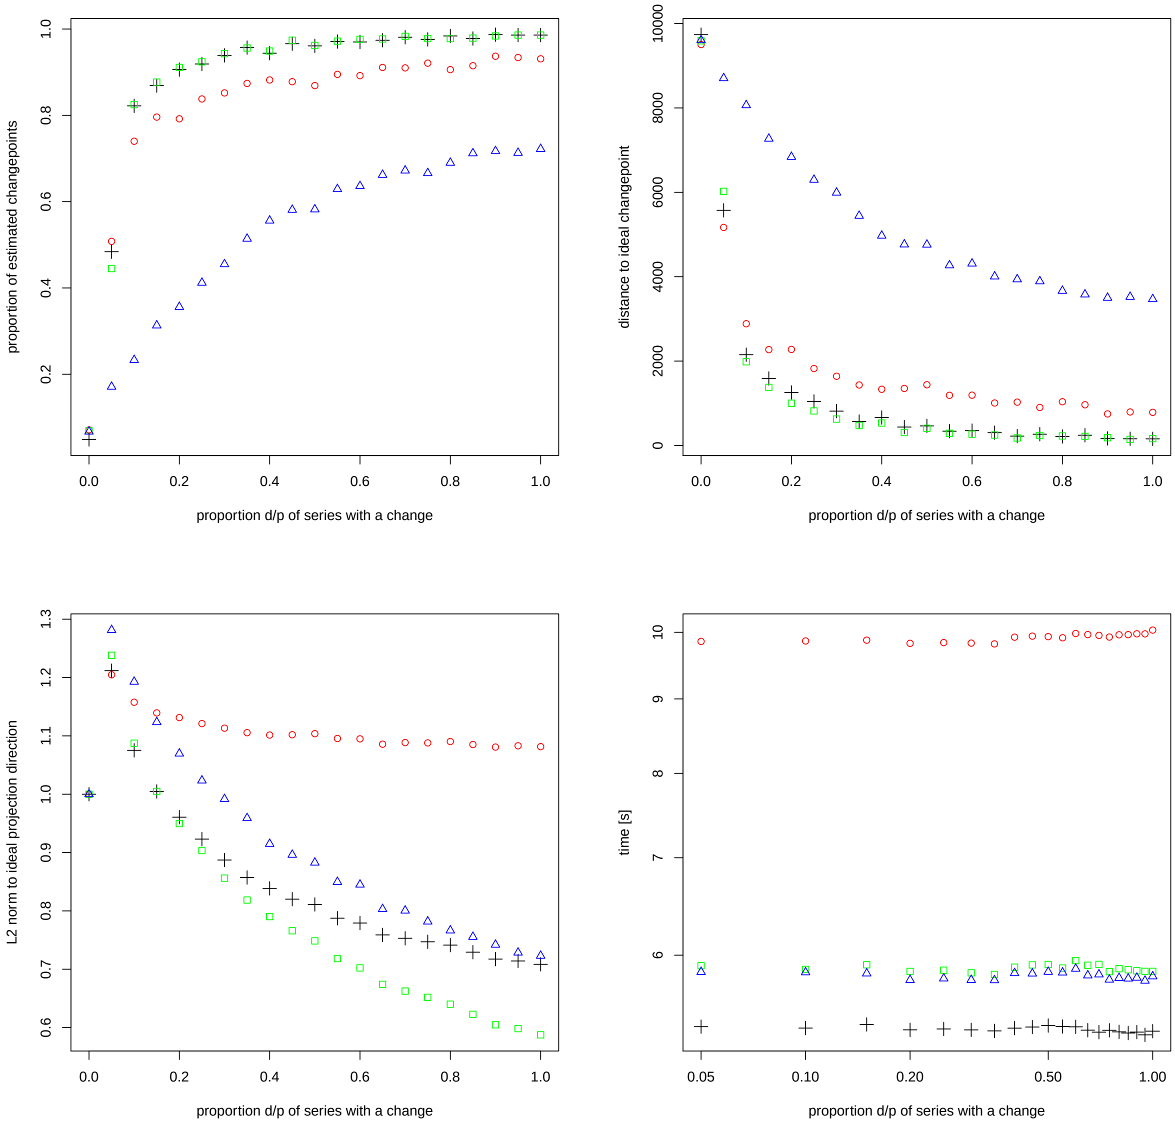

Figure 3: Dependence of BayesProject (black crosses), Inspect (red circles), sum-cusum (green squares), and max-cusum (blue triangles) on the proportion of series d/p with a change while $n = 10000$, $p = 1000$, $s = 0.04$. Plot shows proportion of estimated changes (top left), average distance to the true changepoint location (top right), L_2 distance between the estimated and ideal projection directions (bottom left) and runtime as log-log plot (bottom right).

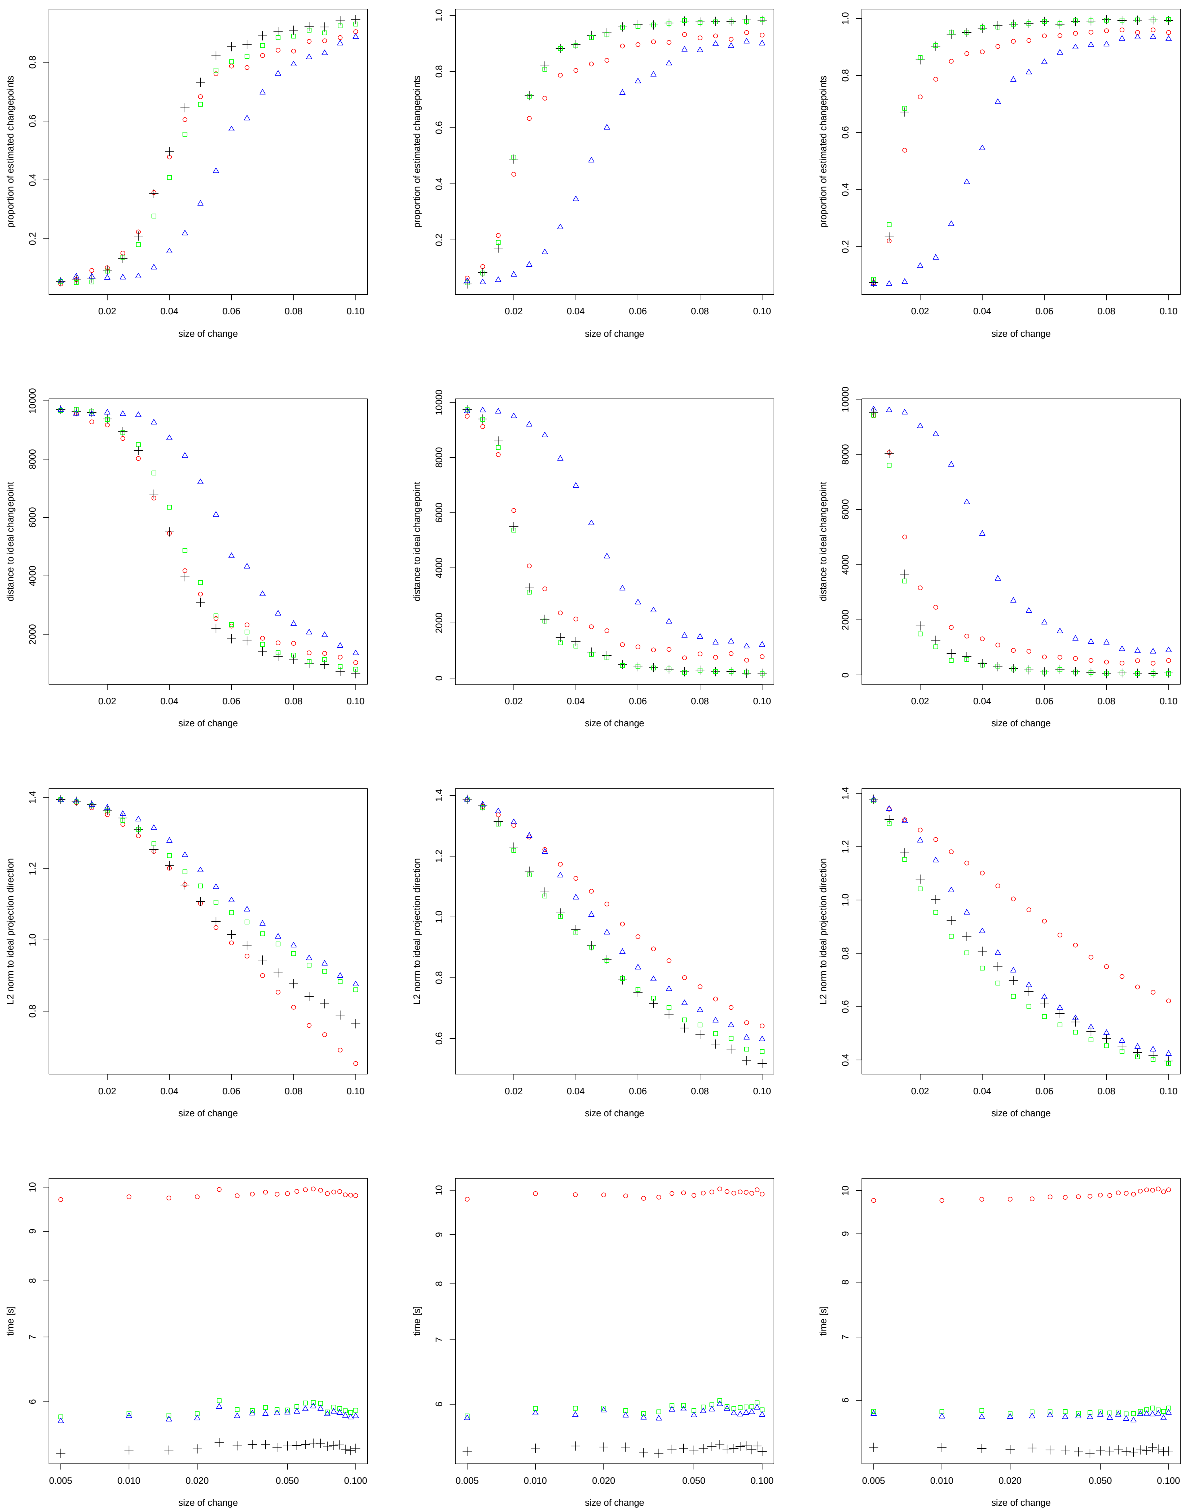

Figure 4: Dependence of BayesProject (black crosses), Inspect (red circles), sum-cusum (green squares), and max-cusum (blue triangles) on the size of the change s while $n = 10000$, $p = 1000$. Rows show proportion of estimated changes (first row), average distance to the true changepoint location (second row), L_2 distance between the estimated and ideal projection directions (third row), and runtime (fourth row). Columns show sparse (left), moderate (middle), and dense (right) scenarios. Log-log plots in the fourth row.

the other two in terms of the average distance to the true changepoint location. Regarding the accuracy of the estimated projection direction, Inspect is most accurate for sparse scenarios and loses accuracy for moderate and denser scenarios, whereas sum-cusum and BayesProject are most accurate for moderate and denser scenarios. The projection direction estimated by max-cusum becomes more accurate than the one of Inspect for denser scenarios. The empirical runtime analysis of all methods shows no dependence on s.

Dependence on the changepoint location. Figure 5 shows the dependence of the four methods on the location of the changepoint. Figure 5 shows that Inspect performs best in terms of the average distance to the true changepoint location in the sparse scenario. BayesProject and sum-cusum become better as the changepoint scenario becomes more dense. The same observation holds true regarding the accuracy of the estimated projection direction. The empirical runtime analysis of all methods shows no dependence on the changepoint location.

2 Dependence on the grid spacing parameter γ and the ratio of ω/σ in eq. (5) of the main article

Choice of γ. Figure 6 investigates the dependence of the accuracy of the estimated projection direction on the choice of γ. The parameter γ determines the grid on which the estimate (6) is computed (Section 3.2). The figure is computed for a scenario with $n = 1000$, $p = 100$, $d = 10$, and a change size of $s = 0.1$. For each choice of γ, the corresponding data point in Figure 6 is the average difference (measured in L_2 distance) among 1000 repetitions between the ideal projection direction and the estimated direction of BayesProject which is computed as in Algorithm 1 of the main article. The figure shows that, as expected, the accuracy of the estimated projection direction increases with the size of the grid (i.e., a higher value of γ). Nevertheless, the narrow range of values depicted on the y-axis suggests that the method is quite robust against the precise choice of γ. In particular, the accuracy of BayesProject does not seem to depend too much on the precise choice of γ as long as the grid is not too sparse (i.e., γ is not too small). Since BayesProject becomes slower with an increasing density of the grid, we fixed $\gamma = 0.6$ in all simulations (see Section 4 of the main article). This is essentially an arbitrary choice.

Dependence on σ^2/ω^2. In eq. (5) of the main article, the assumption $\sigma^2/\omega^2 \approx 0$ is made in order to arrive at the simplified eq. (6). This approximation is motivated by the fact that in a detectable changepoint scenario, we anticipate that the size of the change ω is relatively large compared to the noise in the cusum statistic σ. However, BayesProject seems to be relatively robust against a misspecification of this ratio as demonstrated in Figure 7. In Figure 7, we repeat a simulation study on the size of the change s with three versions of BayesProject. The three versions utilise $\sigma^2/\omega^2 = 0$, $\sigma^2/\omega^2 = 1$ and $\sigma^2/\omega^2 = 2$ in eq. (5). Figure 7 shows that the accuracy of BayesProject seems to be robust against reasonable values of σ^2/ω^2.

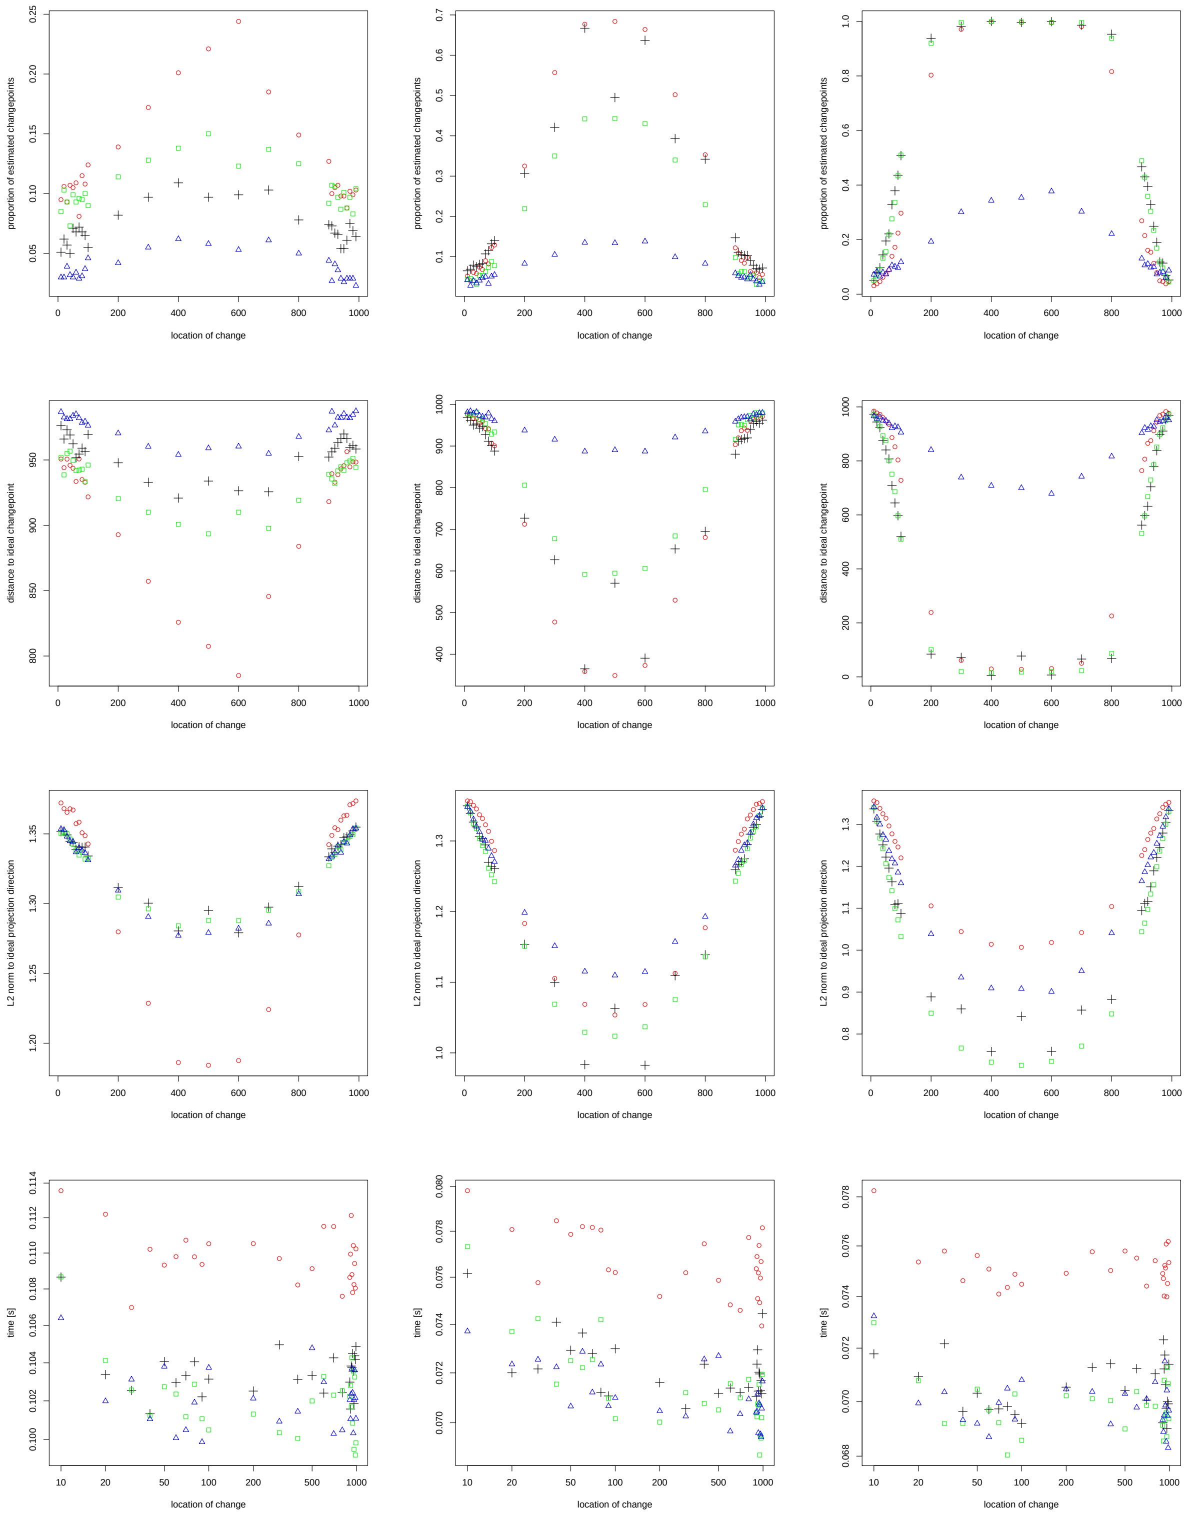

Figure 5: Dependence of BayesProject (black crosses), Inspect (red circles), sum-cusum (green squares), and max-cusum (blue triangles) on the location of the changepoint while $n = 1000$, $p = 100$, $s = 0.1$. Rows show proportion of estimated changes (first row), average distance to the true changepoint location (second row), L_2 distance between the estimated and ideal projection directions (third row), and runtime (fourth row). Columns show sparse (left), moderate (middle), and dense (right) scenarios. Log-log plots in the fourth row.

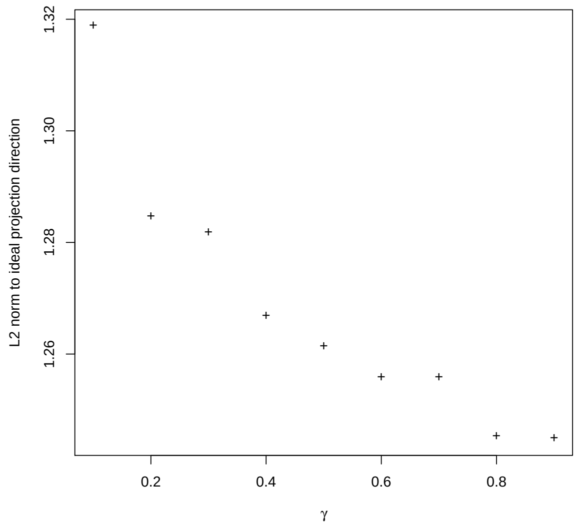

Figure 6: Accuracy of ideal and estimated projection directions (measured in L_2 distance) as a function of the parameter $\gamma \in \{0.1, \ldots, 0.9\}$. Changepoint scenario with $n = 1000$, $p = 100$, $d = 10$, and a change size of $s = 0.1$.

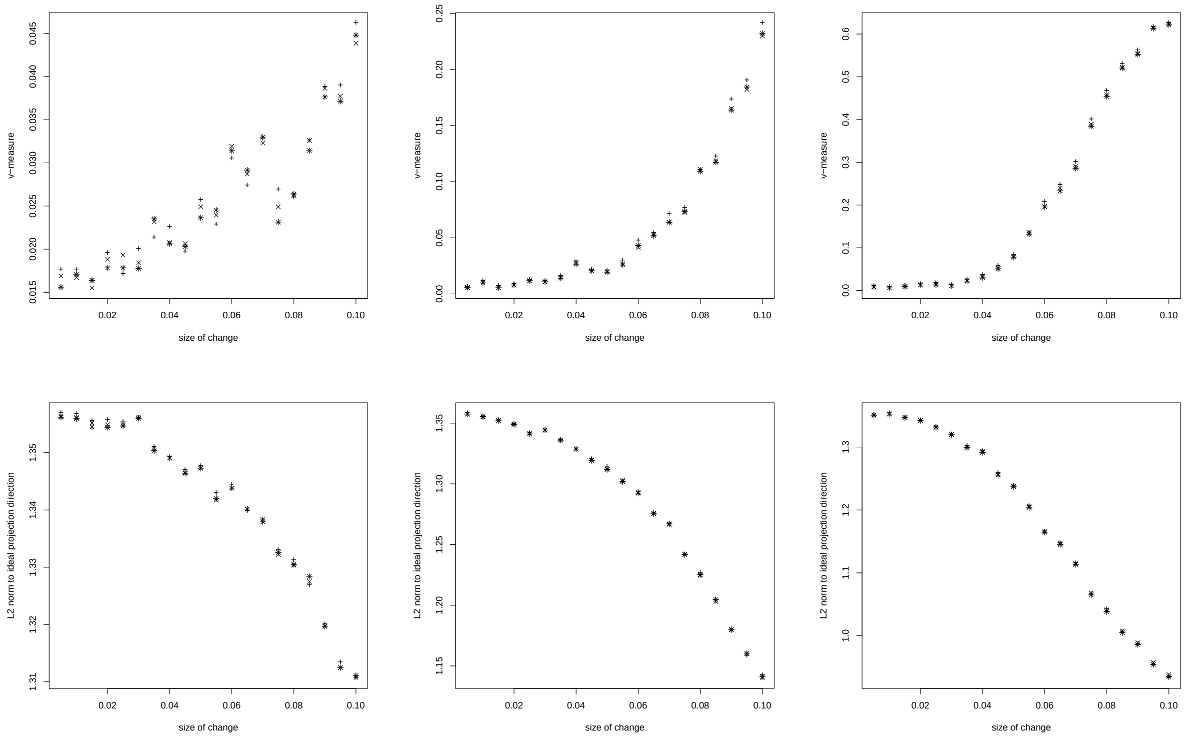

Figure 7: Dependence of BayesProject on the value of σ^2/ω^2 in eq. (5) of the main article for $n = 1000$ and $p = 100$. Plots show dependence on the size of the change s while $\sigma^2/\omega^2 = 0$ (+), $\sigma^2/\omega^2 = 1$ ($\times$) and $\sigma^2/\omega^2 = 2$ ($*$). Rows show accuracy of the found changepoint determined via v-measure (first row) and L_2 distance between the estimated and ideal projection directions (second row). Columns show sparse (left, $d = 5$), moderate (middle, $d = 20$), and dense (right, $d = 50$) scenarios.

3 Dependence on the tuning parameter K

This section repeats some parts of the simulations of both Section 1, as well as of Section 4 of the main article, for three choices of the tuning parameter K (see Section 2.3 of the main article). In particular, those are the simulations which highlight the dependence of the four methods (BayesProject, Inspect, sum-cusum, and max-cusum) on the series length n, the number of series p, the proportion d/p of series with a change, and the size s of the change. The three values of K are: the oracle value of K computed with eq. (9) (see the main article) under the assumption that the true proportion $\bar{\pi}$ and the true quantity ω/σ are known, as well as an overestimate $2K$ and an underestimate $K/2$.

In constrast to the simulations in the previous section and the main article, we decrease the size of the change with the density of the changepoint setting: sparse ($d/p = 0.05$ and $s = 0.09$), moderate ($d/p = 0.2$ and $s = 0.06$), and dense ($d/p = 0.5$ and $s = 0.02$). The size of the change s was selected for each scenario in such a way as to guarantee a proportion of estimated changepoints of roughly 0.8 across all methods.

Figures 8 to 11 show the performance of BayesProject for the three values K, $K/2$ and $2K$. As quality measures, the plots display the proportion of times a timepoint is marked as a changepoint (the proportion of estimated changepoints), the average distance to the true changepoint location, and the accuracy of the estimated projection direction measured in L_2 distance to the ideal projection direction. All results in this section are based on 100 repetitions.

Dependence on n. Figure 8 shows the dependence of BayesProject on the series length n for the three choices of K while $p = 100$ is fixed. It is seen that BayesProject performs very similarly for all choices of K. Regarding the accuracy of the estimated projection direction, the underestimate $K/2$ seems to yield estimates with a lower L_2 distance to the ideal projection direction in sparse scenarios, however at the expense of being least accurate in dense scenarios, and conversely for the overestimate $2K$.

Dependence on p. Figure 9 shows the dependence of BayesProject on the number of series p for the three choices of K while $n = 10000$ is fixed. Similarly to the dependence study on n, Figure 8 shows that BayesProject performs very similarly for all choices of K. As noted for Figure 8, the projection direction estimate is slightly better for $K/2$ or $2K$ in some scenarios, however at the expense of performing worse in the other scenarios.

Dependence on d/p. Figure 10 shows the dependence of BayesProject on the proportion of series d/p with a change for the three choices of K while $n = 10000$, $p = 1000$, and $s = 0.06$ are fixed (where $d < p$ is the number of series with a change). As can be seen, the performance of BayesProject is again almost identical for all choices of K. In the right plot, the same observations as in Figures 8 and 9 regarding the estimated projection direction apply: For sparse scenarios, BayesProject seems to be slightly more accurate when employed with an underestimate $K/2$. As the changepoint scenario becomes denser, BayesProject with the overestimate $2K$ becomes better.

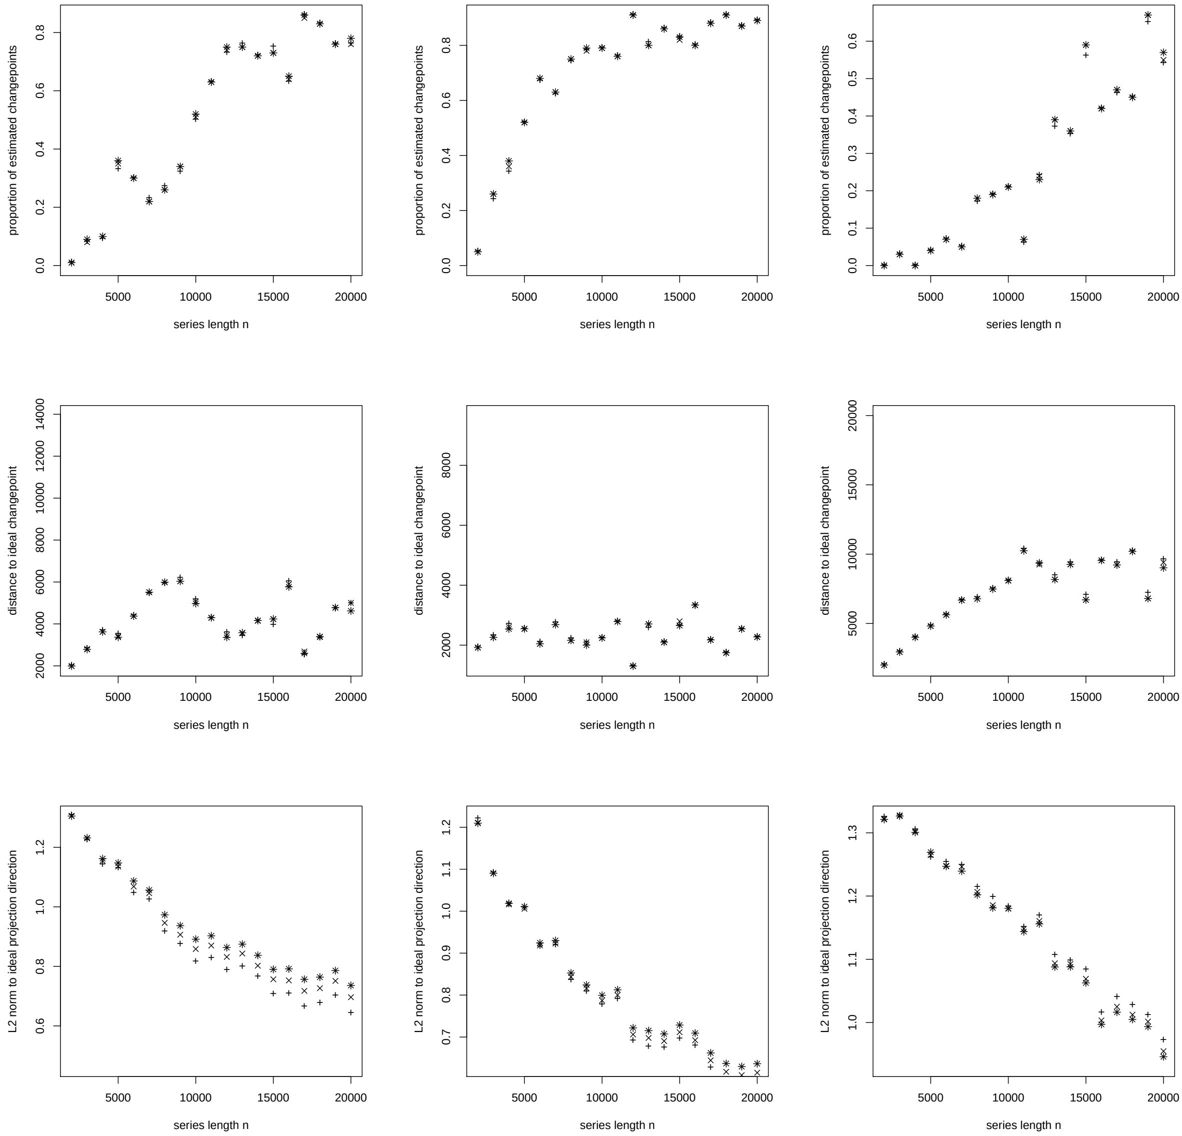

Figure 8: Dependence of BayesProject on the series length n while $p = 100$. Rows show proportion of estimated changes (first row), average distance to the true changepoint location (second row), and L_2 distance between the estimated and ideal projection directions (third row). Columns show sparse (left), moderate (middle), and dense (right) scenarios. BayesProject with optimal K ($\times$), with $K/2$ ($+$) and $2K$ ($*$).

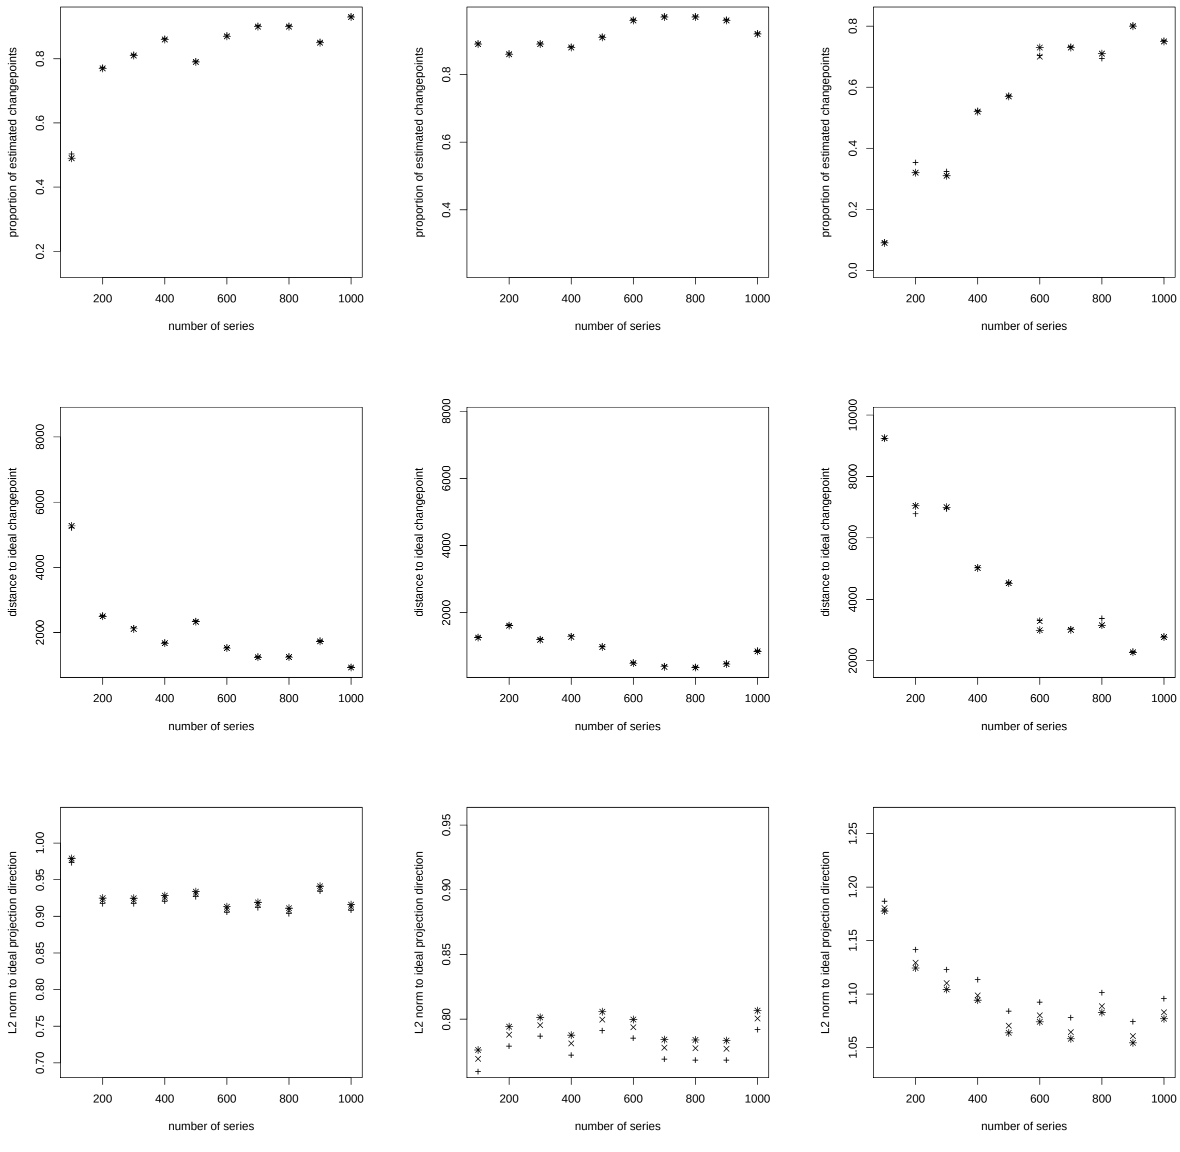

Figure 9: Dependence of BayesProject on the number of series p while $n = 10000$. Rows show proportion of estimated changes (first row), average distance to the true changepoint location (second row), and L_2 distance between the estimated and ideal projection directions (third row). Columns show sparse (left), moderate (middle), and dense (right) scenarios. BayesProject with optimal K ($\times$), with $K/2$ ($+$) and $2K$ ($*$).

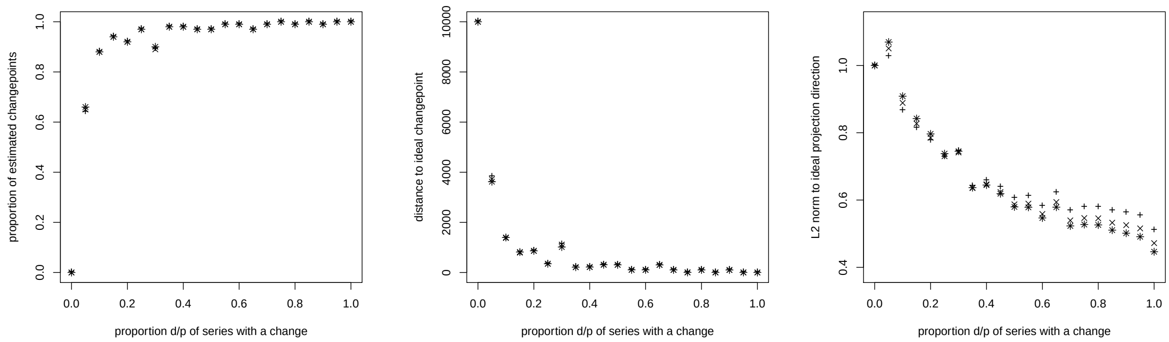

Figure 10: Dependence of BayesProject on the proportion of series d/p with a change while $n = 10000$, $p = 1000$ and $s = 0.06$. Plot shows the proportion of estimated changes (left), the average distance to the true changepoint location (middle), and the L_2 distance between the estimated and ideal projection directions (right). BayesProject with optimal K ($\times$), with $K/2$ ($+$) and $2K$ ($*$).

Dependence on s. Figure 11 shows the dependence of BayesProject on the size s of the change for the three choices of K while $n = 10000$ and $p = 1000$ are fixed. As can be seen, the performance of BayesProject (under all measures) is again almost identical for the three choices of K.

4 Univariate case

Figure 12 shows results for univariate changepoint detection. We observe that all methods perform roughly equally, and that results are qualitatively similar to the multivariate case.

5 Null case

Figure 13 shows the proportion of detected changes for all four methods in the null case in which no changepoint is present. The threshold for all methods was computed as described in the main article as the 95% quantile of the maximum cusum statistic on normal data without a changepoint. Figure 13 shows that all methods roughly keep the 5% rate, with possibly a slight overestimation of the presence of changepoints. We observe that the proportion of erroneously found changepoints is roughly the same for all methods, in particular BayesProject and Inspect perform similarly.

6 Asymptotic runtimes of Inspect with SVD and partial SVD (Lanczos iterations)

For their Inspect algorithm, Wang and Samworth (2017) compute the left singular vector using an SVD calculation, see e.g. the *InspectChangepoint* package on CRAN (Wang and Samworth, 2016) which calls the R function *eigen* with parameter *symmetric=true*. However, a potentially

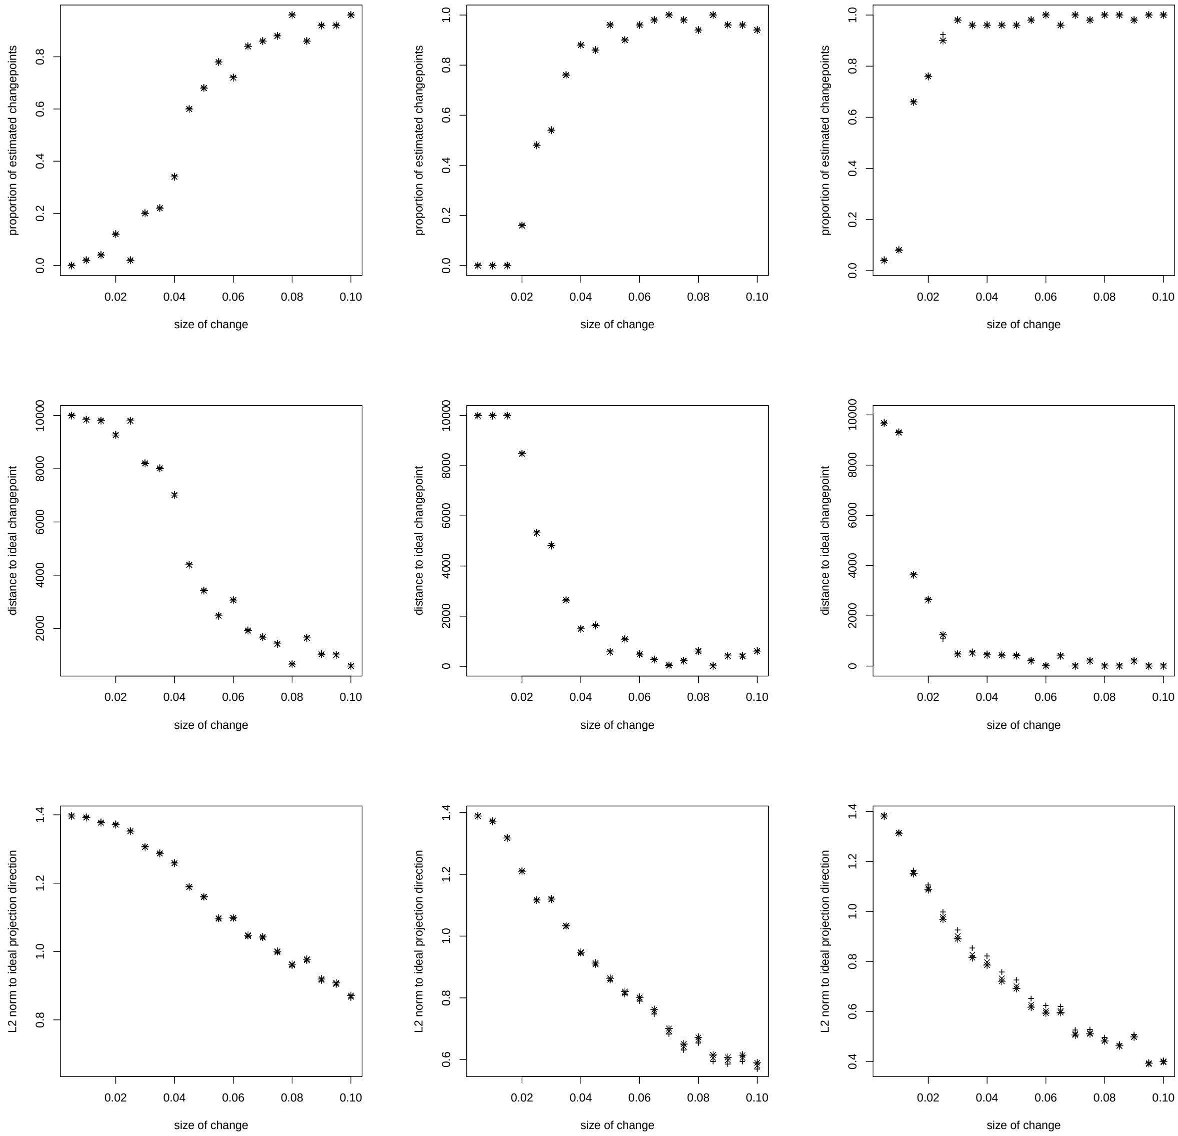

Figure 11: Dependence of BayesProject on the size of the change s while $n = 10000$ and $p = 1000$. Rows show proportion of estimated changes (first row), average distance to the true changepoint (second row), and L_2 distance between the estimated and ideal projection directions (third row). Columns show sparse (left), moderate (middle), and dense (right) scenarios. BayesProject with optimal K ($\times$), with $K/2$ ($+$) and $2K$ ($*$).

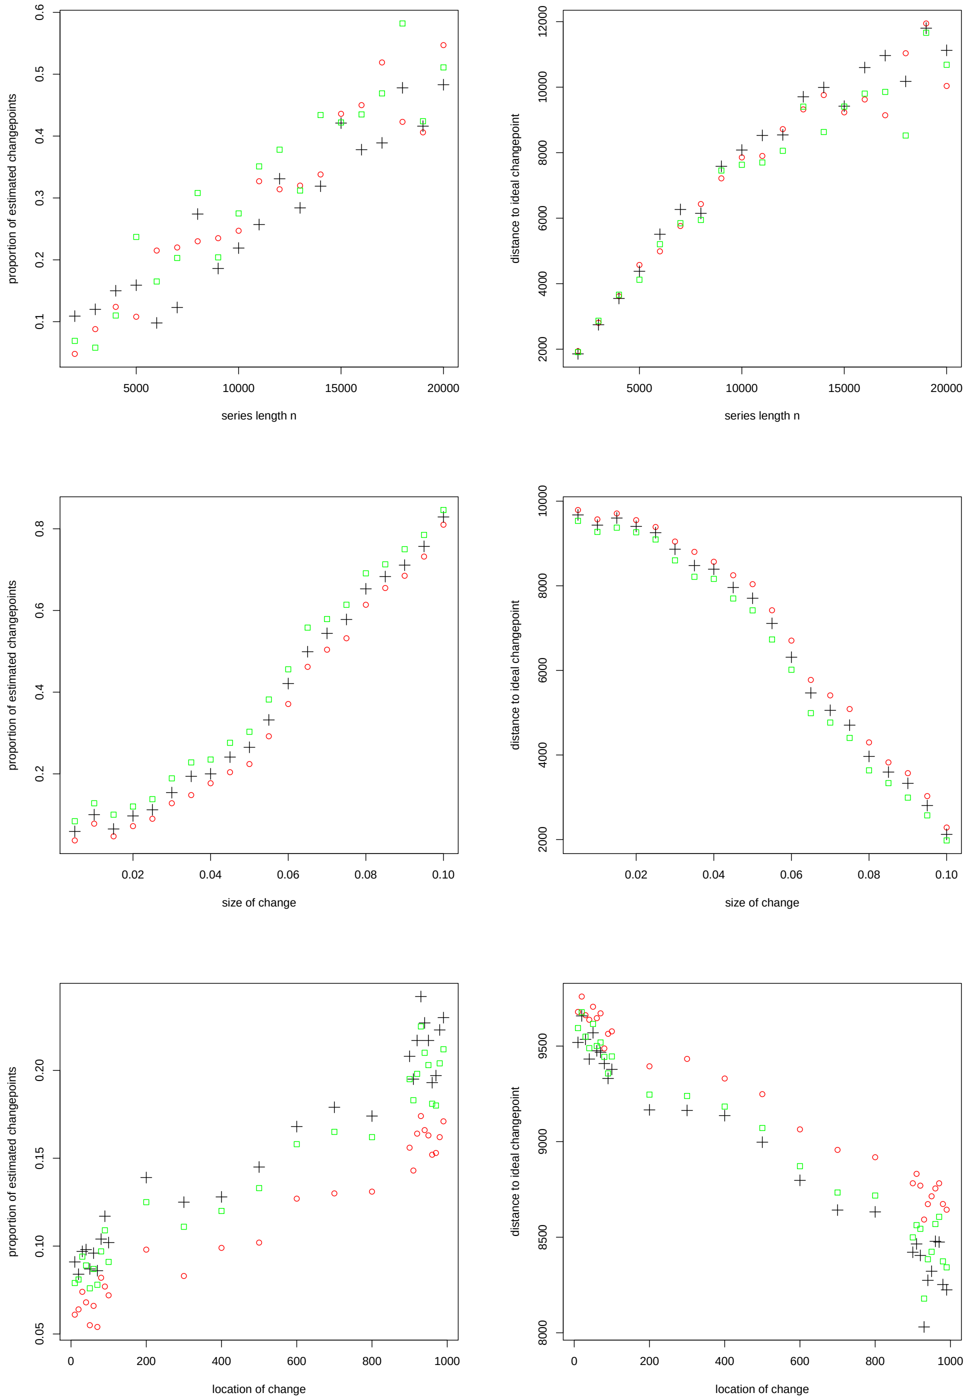

Figure 12: Univariate changepoint detection (case $p = 1$). Dependence of BayesProject (black), Inspect (red), sum-cusum (green), and max-cusum (blue) on the series length n (first row), the size of the change s (second row), and the location of the changepoint (third row). In the first row we let $n \in \{200, \ldots, 20000\}$, while we fix $n = 10000$ in the second and third rows. In the second row, the size of the change s varies, while $s = 0.05$ is fixed in the first and third rows. Columns show proportion of estimated changepoints (left) and average distance to the true changepoint location (right).

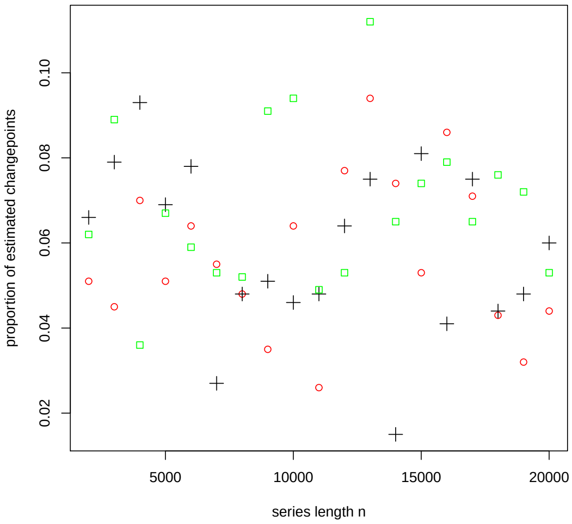

Figure 13: No changepoint present in the data. Proportion of detected changes for BayesProject (black), Inspect (red), sum-cusum (green), and max-cusum (blue) as a function of the series length n while $p = 100$.

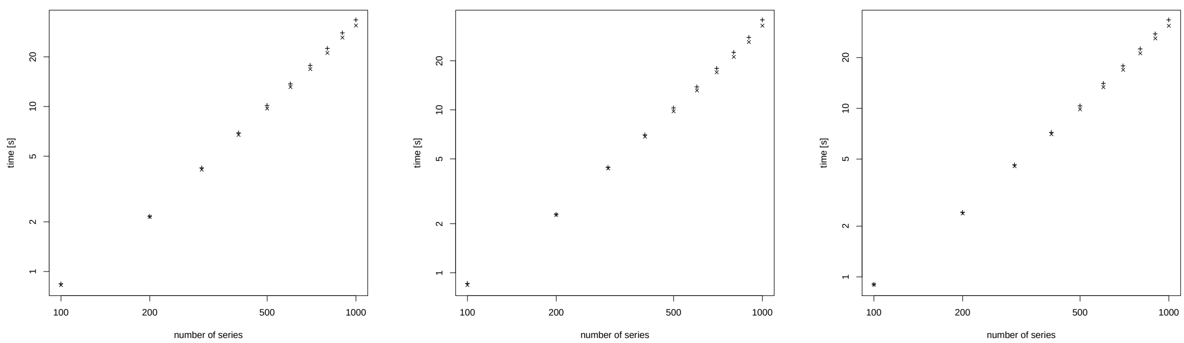

Figure 14: Runtimes of Inspect when computing the leading left singular vector via SVD (+) and Lanczos iterations ($\times$) as a function of p while $n = 10000$. Plots show runtimes in seconds for the sparse (left, $d/p = 0.05$), moderate (middle, $d/p = 0.2$), and dense (right, $d/p = 0.5$) scenarios. Size of change $s = 0.05$. Slope estimates for the moderate scenario (middle) are 1.62 (SVD, (+)) and 1.59 (Lanczos, ($\times$)).

more efficient way is to use a partial SVD (Lanczos iterations). The R-package *RSpectra* (Qiu et al., 2019) provides such a Lanczos implementation.

We modify Inspect to compute the leading left singular vector via a partial SVD computation in its function *sparse.svd*. For this we employ the function *eigs_sym* from the RSpectra package (since it is known that the input matrix is symmetric). As the function *eigs_sym* is not available in Rcpp, we do not employ the C++ implementation of Inspect we used in all other experiments in the main paper and the supplementary material, but rather use plain R code for all experiments as in the original InspectChangepoint package. We are interested in the asymptotic runtime scalings of the Inspect algorithm. Since the SVD is computed on $\hat{M}\hat{M}^\top \in \mathbb{R}^{p\times p}$ or $\tilde{M}\tilde{M}^\top \in \mathbb{R}^{p\times p}$ depending on whether soft-thresholding is being used (see Section 2.2 of the main article), both of which depend on p only, it suffices to consider the scaling in p.

Figure 14 shows runtimes as a function of p while $n = 10000$ as in the simulations of main article. We observe that using a partial SVD leads to a slightly improved asymptotic runtime of Inspect compared to the standard SVD approach.

7 Density plots

This section provides complete density plots for all scenarios considered in Tables 1, 2, and 3 in the main article.

Figures 15 and 16 show density plots for all scenarios given in Table 1 in the main article (the original setting adopted from Wang and Samworth (2017)). Figure 15 displays the case $\vartheta = (0.6, 1.2, 1.8)$, and Figure 16 displays the case $\vartheta = (0.4, 0.8, 1.2)$. We see that BayesProject does not perform very well for the first changepoint, while Inspect, Enikeeva and Harchaoui (2019), and Yu and Chen (2020) perform very well. For the second and third changepoints, Enikeeva and Harchaoui (2019) and BayesProject become more competitive.

Figure 17 shows density plots for all scenarios in Table 2 of the main article (having more uneven segment lengths). We see that in the complete overlap scenario, BayesProject is best for the first two changepoints but does get outperformed for the third changepoint. For the half and no overlap scenario, Enikeeva and Harchaoui (2019) as well as Yu and Chen (2020) are best for the left changepoint, while Inspect and BayesProject catch up for the second and third changepoint.

Figure 18 shows density plots for all scenarios in Table 3 of the main article. We see that for two changepoints on the sides, BayesProject performs among the best algorithms in all three scenarios and for both changepoints, though challenged by Yu and Chen (2020) for the first changepoint.

References

Cho, H. (2016). Change-point detection in panel data via double CUSUM statistic. *Electronic Journal of Statistics*, 10:2000–2038.

Enikeeva, F. and Harchaoui, Z. (2019). High-dimensional change-point detection under sparse alternatives. *Ann Stat*, 47(4):2051–2079.

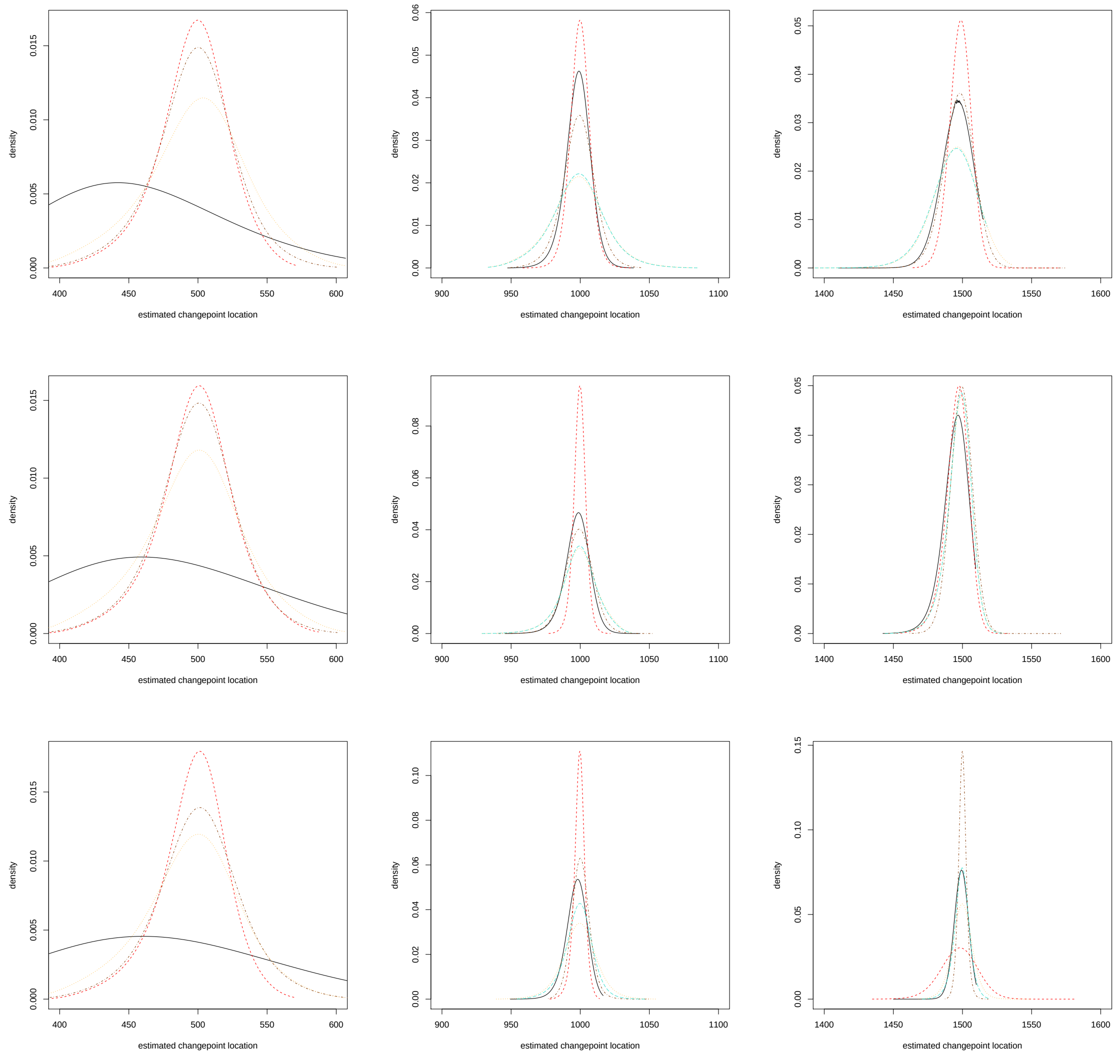

Figure 15: Estimated densities of the location of changepoint estimates reported in Table 1 for $\vartheta = (0.6, 1.2, 1.8)$ of the main article for the complete (top row), half (middle row) and no overlap (bottom row) scenario. BayesProject (black, solid line), Inspect (red, dashed line), Cho (2016) (orange, dotted line), Enikeeva and Harchaoui (2019) (brown, dotdashed line), Yu and Chen (2020) (turquoise, longdashed line). Changepoints 500 (left), 1000 (middle) and 1500 (right).

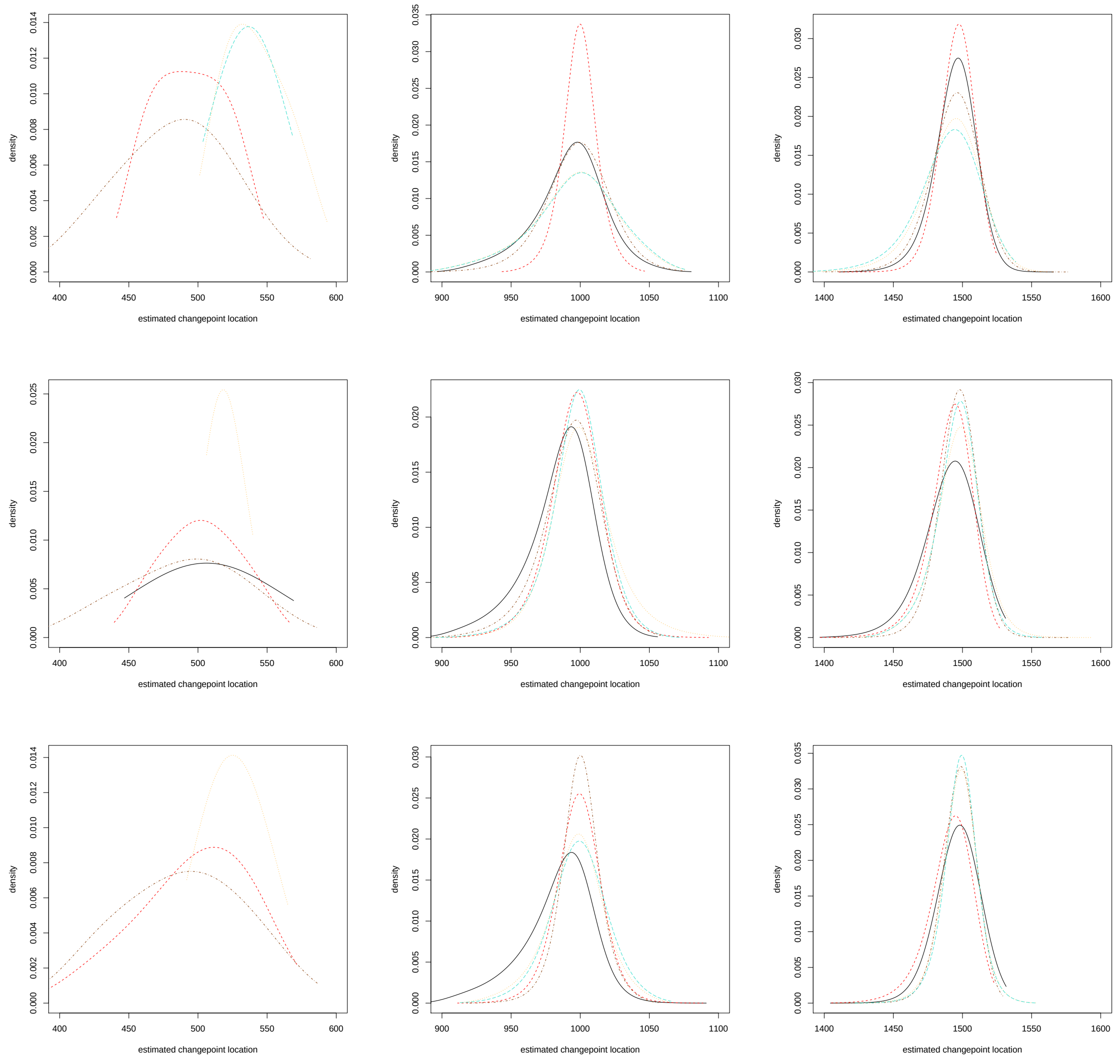

Figure 16: Estimated densities of the location of changepoint estimates reported in Table 1 for $\vartheta = (0.4, 0.8, 1.2)$ of the main article for the complete (top row), half (middle row) and no overlap (bottom row) scenario. BayesProject (black, solid line), Inspect (red, dashed line), Cho (2016) (orange, dotted line), Enikeeva and Harchaoui (2019) (brown, dotdashed line), Yu and Chen (2020) (turquoise, longdashed line). Changepoints 500 (left), 1000 (middle) and 1500 (right).

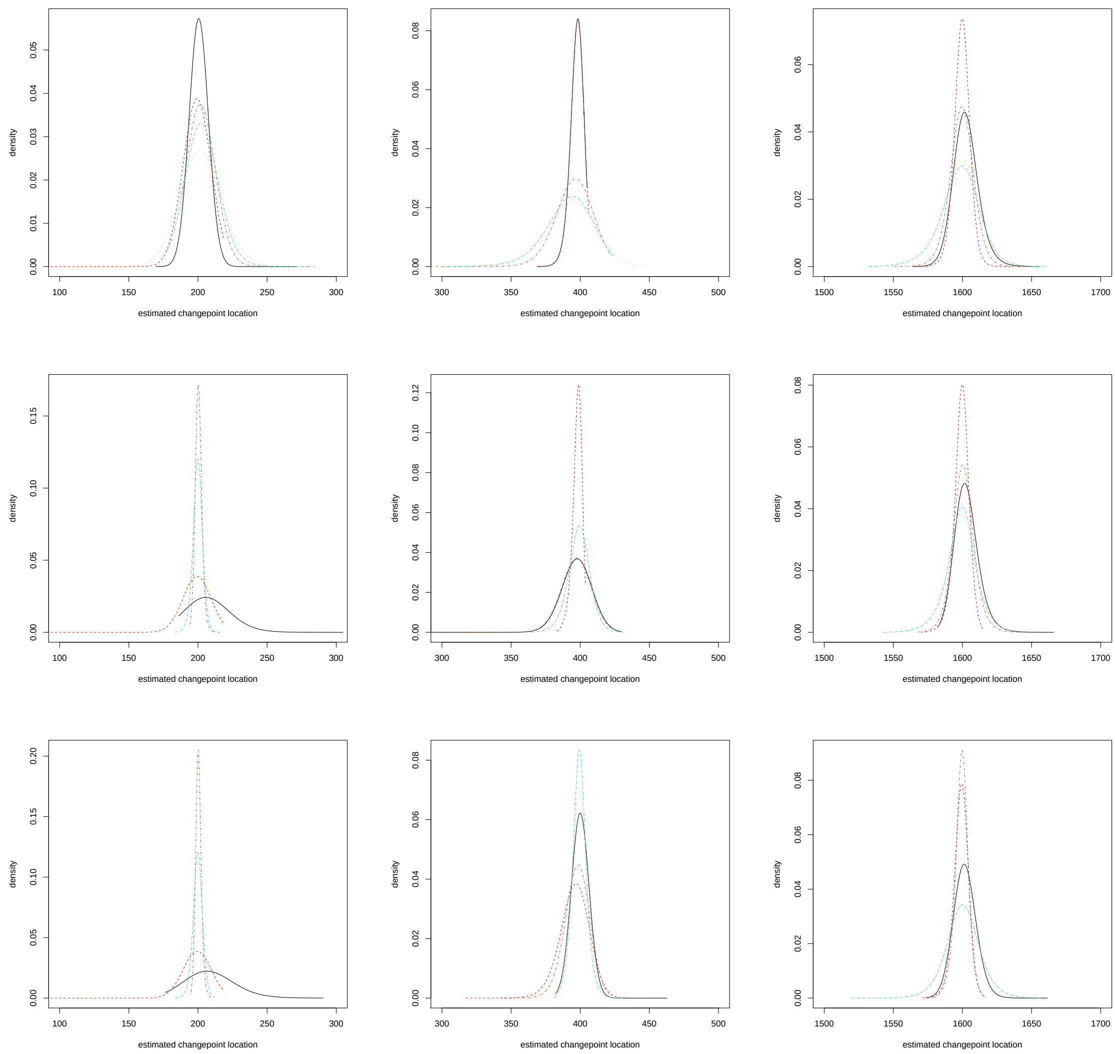

Figure 17: Estimated densities of the location of changepoint estimates reported in Table 2 of the main article for the complete (top row), half (middle row) and no overlap (bottom row) scenario. BayesProject (black, solid line), Inspect (red, dashed line), Cho (2016) (orange, dotted line), Enikeeva and Harchaoui (2019) (brown, dotdashed line), Yu and Chen (2020) (turquoise, longdashed line). Changepoints 200 (left), 400 (middle) and 1600 (right).

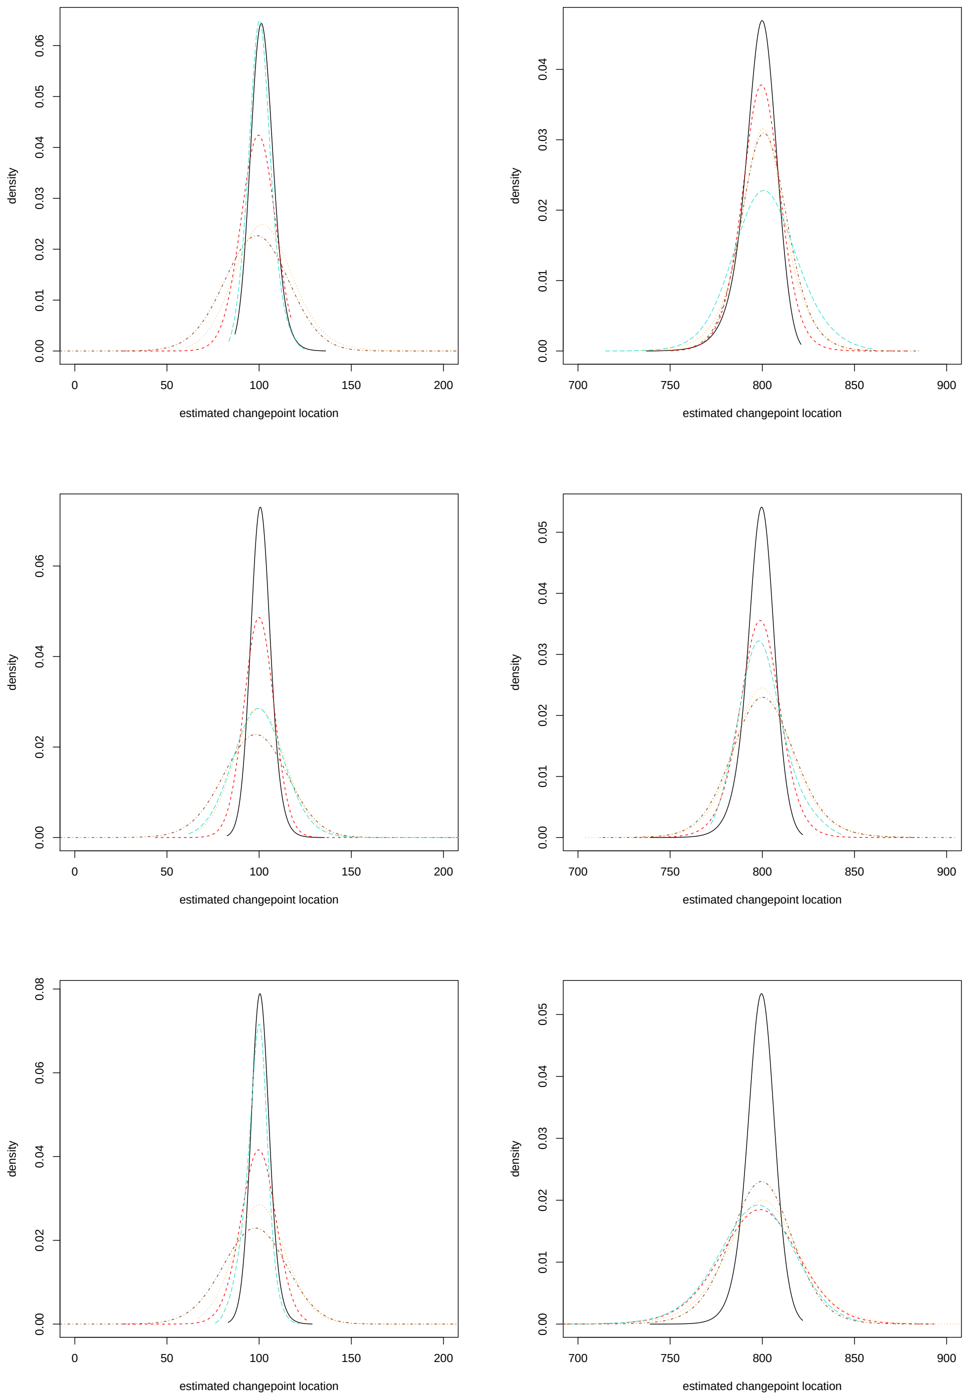

Figure 18: Estimated densities of the location of changepoint estimates reported in Table 3 of the main article for the complete (top row), half (middle row) and no overlap (bottom row) scenario. BayesProject (black, solid line), Inspect (red, dashed line), Cho (2016) (orange, dotted line), Enikeeva and Harchaoui (2019) (brown, dotdashed line), Yu and Chen (2020) (turquoise, longdashed line). Changepoints 100 (left) and 800 (right).

Qiu, Y., Mei, J., Guennebaud, G., and Niesen, J. (2019). RSpectra: Solvers for Large-Scale Eigenvalue and SVD Problems. R package version 0.16-0.

Wang, T. and Samworth, R. (2016). InspectChangepoint: High-Dimensional Changepoint Estimation via Sparse Projection. R package version 1.0.1.

Wang, T. and Samworth, R. (2017). High dimensional change point estimation via sparse projection. *J Roy Stat Soc B Met*, 80(1):57–83.

Yu, M. and Chen, X. (2020). Finite sample change point inference and identification for high-dimensional mean vectors. *arXiv:1711.08747*, pages 1–71.