



Detection of spatiotemporal changepoints: a generalised additive model approach

Michael J. Hollaway¹ · Rebecca Killick²

Received: 19 December 2023 / Accepted: 22 July 2024
© The Author(s) 2024

Abstract

The detection of changepoints in spatio-temporal datasets has been receiving increased focus in recent years and is utilised in a wide range of fields. With temporal data observed at different spatial locations, the current approach is typically to use univariate changepoint methods in a marginal sense with the detected changepoint being representative of a single location only. We present a spatio-temporal changepoint method that utilises a generalised additive model (GAM) dependent on the 2D spatial location and the observation time to account for the underlying spatio-temporal process. We use the full likelihood of the GAM in conjunction with the pruned linear exact time (PELT) changepoint search algorithm to detect multiple changepoints across spatial locations in a computationally efficient manner. When compared to a univariate marginal approach our method is shown to perform more efficiently in simulation studies at detecting true changepoints and demonstrates less evidence of overfitting. Furthermore, as the approach explicitly models spatio-temporal dependencies between spatial locations, any changepoints detected are common across the locations. We demonstrate an application of the method to an air quality dataset covering the COVID-19 lockdown in the United Kingdom.

Keywords Changepoint · Spatio-temporal · PELT · GAM

1 Introduction

In time series analysis, a changepoint is a point in time where an abrupt change in the statistical properties of the time series occurs. Typically, the focus of changepoint analysis has been on detecting changes in the temporal structure of the time series, often changes in mean, variance, trend or a combination (Beaulieu and Killick 2018). Changepoint detection is utilised in a wide range of fields including genomics (Caron et al. 2012; Liehrmann et al. 2023), health (Younes et al. 2019; Tapsoba et al. 2020; Creswell et al. 2023), and environmental science (Lund et al. 2007; Gallagher et al. 2012; Beaulieu et al. 2020).

Historically, changepoint detection has focused on the univariate case whereby changes are detected in a single time series although in recent years this has been extended to multivariate time series (Ma and Yau 2016; Hahn et al. 2020; Lowther et al. 2023). The majority of multivariate approaches harness detection power across the series but still make the assumption of independence between them.

However, often with environmental datasets, we are presented with data from various spatial locations that measure the evolution of a variable not only over time but also the spatial domain on which the measurement process lies, thus violating this assumption of independence between time series. Furthermore, univariate methods can only be deployed in a marginal sense and the changes detected are representative of a single spatial location only. This property of environmental datasets has prompted new approaches to be developed to account for the dependence between the multiple time series. Ryan and Killick (2023) detect changes in covariance but focus on the second order and assume the mean structure is zero. Gromenko et al. (2017) use an approach based on functional data analysis in order to detect a single changepoint in the annual pattern of precipitation data at fixed spatial locations. In this case, observations

✉ Michael J. Hollaway
mhollaway@ceh.ac.uk

Rebecca Killick
r.killick@lancaster.ac.uk

¹ UK Centre for Ecology & Hydrology, Lancaster Environment Centre, Library Avenue, Lancaster LA1 4AP, UK

² School of Mathematical Sciences, Lancaster University, Lancaster LA1 4YF, UK

are modelled as functional valued time sequences across the multiple spatial locations with the addition of spatially correlated error functions. Dette and Quanz (2023) take spatio-temporal changepoint detection further by focusing on changes exceeding a certain threshold rather than considering the problem of arbitrary change sizes. Finally, Zhao et al. (2024) develop a composite likelihood approach using a piecewise stationary spatio-temporal process in order to detect underlying changes in the non-stationary spatio-temporal process across multiple spatial locations. By using a pairwise composite likelihood in conjunction with the Pruned Linear Exact Time (PELT; Killick et al. (2012)) algorithm, they are able to overcome the computational burden of spatio-temporal modelling to detect multiple changepoints.

There are many spatio-temporal process models that have been developed in the literature coming from different starting paradigms. Examples include linear mixed effect models that can model spatial and temporal effects (e.g. using the R package `nlme` based on the approaches of Laird and Ware (1982); Lindstrom and Bates (1988)), Bayesian univariate and multivariate spatio-temporal random effects models (Finley et al. 2015; Finley and Banerjee 2020) and Bayesian hierarchical approaches (Bakar and Sahu 2015). However, one of the most easily accessible spatio-temporal processes for practitioners is the generalised additive model (GAM), due to the conceptual extended regression framework, and availability of code and accessible introductions. In this paper, we adopt a spatio-temporal process using a GAM which is dependent on the 2-D spatial location and the time of the observation. We then use the full likelihood in conjunction with PELT to detect multiple changepoints across spatial locations in a computationally efficient manner.

The remainder of the paper is set out as follows: Sect. 2 provides a description of the derivation of the method and its application to spatio temporal changepoint detection. The approach is demonstrated through series of simulation studies in Sect. 3 and application to a real world example using air quality data over the United Kingdom (UK) in Sect. 4. Finally Sect. 5 presents concluding remarks.

2 Methods

There are two main components needed for detecting change; 1) The model to fit between two changepoints, and 2) the algorithm for identifying changes. The spatiotemporal model we utilise between changepoints is a Generalised Additive Model (GAM) (Wood 2017). The likelihood of the GAM is utilised within the Pruned Exact Linear Time (Killick et al. 2012) algorithm for detecting multiple changepoints over time. These are described in the remainder of this section.

2.1 GAM model

Let $y_{s,t}$ be a three-dimensional observation of a process of interest over a 2-d space, $s = (u, v)$ and time, t . The collection of $\{y_{s,t}\}_{(s \in (U, V), t=1, \dots, n)}$ is a spatio-temporal process over a defined spatial domain (U, V) observed at n time points. This could be, for example, air quality (NO_2 or O_3) observed at different spatial locations across the UK over time. We choose to fit a generalised additive model (GAM) to data of this type.

$$y_{s,t} = f_1(x_s) + f_2(x_t) + f_3(x_s, x_t) + \epsilon_{s,t} \quad (1)$$

where $f_1(\cdot)$ is a function over 2-d space, $f_2(\cdot)$ is a function over time and $f_3(\cdot)$ is a function over both time and space. The $\epsilon_{s,t}$ are errors that are independent of all fitted components with mean 0 and variance σ^2 .

There are many different functional forms that the $f_i(\cdot)$ can take, including thin plate and cubic spline regressions, and tensor products. See Wood (2017) for descriptions. We focus on a GAM described by (1) here but additional explanatory covariates can also be added if warranted.

Recall that in identifying changepoints we seek to identify changes in the GAM model parameters. To do this we need to have a way of describing and comparing the fit of different GAM models to different segments of the data. A commonly used measure of fit is the likelihood and we will adopt this approach, including maximum likelihood estimation for the parameter estimates.

2.2 Changepoint estimation

In describing the GAM model in Sect. 2.1 we sought to optimise the no changepoint scenario

$$\sum_{i=1}^n \mathcal{C}(y_{s,i} | \theta) \quad (2)$$

for θ . Here $\mathcal{C}(\cdot)$ is a measure of fit given fitted parameters $\hat{\theta}$ which is twice the negative log-likelihood of the GAM model (1). As written, equation (2) fits a single GAM model (and parameters via maximum likelihood) for all time points $i = 1, \dots, n$. To add changepoints we focus on the time component in what follows. The spatial component, and the within-segment space-time interactions, are dealt with by the GAM modelling for each segment. Thus we are detecting changes in time of the spatial-temporal process parameters, we do not seek to detect changes across space i.e., cliffs.

Recall that we define changepoints at times $0 = \tau_0, \tau_1, \dots, \tau_M, \tau_{M+1} = n$. Under the changepoint assumption the model parameters are restricted to be the same across segments of data and we seek to optimise,

$$\min_{\tau, M} \sum_{m=0}^M \sum_{i=\tau_m+1}^{\tau_{m+1}} \mathcal{C}(y_{s,i} | \hat{\theta}_m). \tag{3}$$

Due to the discrete nature of both the number and location of changepoints, standard estimation methods cannot be directly applied. We now have a model selection problem where you need to select the appropriate number of changepoints. This is akin to choosing the number of regressors in a regression problem. Without restrictions, the optimisation of (3) would choose the maximum number of changepoints, M and so as in the regression context, we need to penalize. Zheng et al. (2022) demonstrates that penalties of the form $CM \log(n)$ are consistent for likelihood-based cost functions, $\mathcal{C}(\cdot)$, and constant C with respect to n . Thus we optimize,

$$F(n) = \min_{\tau, M} \sum_{m=0}^M \sum_{i=\tau_m+1}^{\tau_{m+1}} \mathcal{C}(y_{s,i} | \hat{\theta}_m) + CM \log(n). \tag{4}$$

Optimizing (4) over all possible combinations of M and τ is a computationally intensive task. Killick et al. (2012) demonstrates how a combination of dynamic programming and pruning the search space can reduce the computational burden from $\mathcal{O}(2^n)$ to $\mathcal{O}(n)$. With the assumed independence of the segments in (4), dynamic programming allows us to rewrite the search for all changepoints in (4) into the search for the last changepoint prior to n ,

$$F(n) = \min_{\tau^*} F(\tau^*) + \sum_{i=\tau^*}^n \mathcal{C}(y_{s,i} | \hat{\theta}_M) + C \log(n). \tag{5}$$

Computing $F(\tau^*)$ recursively for $\tau^* = 1, \dots, n$ recovers the optimal set of changepoints for penalty $CM \log(n)$ in $\mathcal{O}(n^2)$ computational time. To reduce this to $\mathcal{O}(n)$ one can prune the minimisation in (5). As the minimisation is looking for the best last changepoint location at each step of the algorithm, where there has been an obvious changepoint prior to the current step, the best last changepoint is unlikely to be before this obvious changepoint. This intuition is mathematically optimal to prune an individual τ^* from the minimization set, if it satisfies

$$F(\tau^*) + \sum_{i=\tau^*}^t \mathcal{C}(y_{s,i} | \hat{\theta}_M) \geq F(t). \tag{6}$$

Intuitively this says that if at any time in the recursive computation, a candidate last changepoint location is more than $C \log(n)$ larger than the optimal likelihood at that time, it can never be the last changepoint at any future point so it can be pruned from the minimisation in (5). The authors call this algorithm, PELT, Pruned Exact Linear Time.

We use PELT as a wrapper for our GAM model by using the negative twice the log-likelihood as $\mathcal{C}(\cdot)$ in (5). We denote this GAM-PELT in the remainder of the paper. The computational cost of this is then $\mathcal{O}(Ln)$ where L is the computational order of evaluating the likelihood for the GAM model in a single segment.

3 Simulations

In this section we evaluate the GAM-PELT method, with default SIC/BIC penalty, to see if it can accurately detect different types of change. In applications, several different types of change can occur so we run several scenarios in which changepoints are specified in the; spatial structure, temporal structure, both (spatio-temporal), or no change at all. It is important to include simulations with no change to ensure that false changepoints are not detected when no changepoints are present. A summary of the scenarios run can be found in Table 1 (no changes) and Table 3 (changes).

The GAM used in our simulations is defined following the form of Eq. 1 where f_1 is a 2D thin plate regression spline over U, V , f_2 is a cubic regression spline over T and f_3 is a tensor product interaction to account for the interactions between the spatial and temporal components. The splines were defined using the default settings from the `mgcv` package with the exception of the number of knots in the cubic regression spline which were set to 5. Naturally other GAM forms could be used depending on the dynamics of a given application.

We compare the performance of the GAM-PELT method with the closest available marginal (univariate) method; the change in mean model with autoregressive errors of order one, AR(1). The marginal method ignores the spatial component and fits each spatial location independently, identifying multiple changepoints with the same PELT search algorithm. Code for this method is available in the `EnvCpt` R package on CRAN (Killick et al. 2021). Both the GAM-PELT and the marginal approach used the standard Bayesian Information Criterion (BIC) as the penalty value in PELT.

For all scenarios the number of time points and spatial locations were fixed at 200 and 50 respectively, and 3 changepoints at timesteps 50, 100 and 150. For each scenario there are 100 replicates with the spatial locations generated at random at the start of each replicate uniformly from $u \sim \text{Unif}(-3,3)$, $v \sim \text{Unif}(40,60)$, rounded to 1 decimal place. To compare the accuracy of the GAM-PELT method and the traditional marginal approach, we consider the timing of the detected changepoints. A changepoint is considered to be accurately detected (i.e. true positive) if it sits within 10 timesteps of the true position. If more than one changepoint sits within this window, one is counted as the true changepoint, and the other as false. Finally, the number of

Table 1 Summary of scenarios where there are no spatial and/or temporal changes in the simulated dataset

Scenario	Time	Space
A	AR1: same all locations	Constant
B	AR1: random all locations	Constant
C	Independent	Constant
D	Independent	Independent
E	Independent	Structured correlation
F	Independent	2-D <i>gamSim</i> example (<i>mgcv</i>)

Table 2 Percentage of estimated changepoints m among 100 replications at 50 locations under various no changepoint scenarios

Scenario	GAM-PELT			Marginal		
	0	1–4	>5	0	1–4	>5
A	98	2	0	98.58	1.42	0
B	92	8	0	98.36	1.64	0
C	100	0	0	99.10	0.90	0
D	100	0	0	99.04	0.96	0
E	100	0	0	99.38	0.62	0
F	100	0	0	99.04	0.96	0

false changepoints (i.e. false positives) is the total number of changepoints minus the number correctly identified. To be fair in the comparison with the marginal approach, we perform this evaluation independently across all spatial locations according to the expected changepoints for each method, and then average across locations. Thus a falsely detected changepoint in GAM-PELT will be counted as 50 false changepoints and a true detection as 50 true changepoints. Conversely, for the marginal approach, where only a single spatial location has a change, if any other location detects a change then it is considered a false changepoint.

3.1 No changes

To ensure that the GAM-PELT method doesn't falsely detect changepoints, we run a series of scenarios that have different spatial and temporal structures but have no changepoints. Table 1 shows a summary of the scenarios run with specific parameter values given in the Supplementary Material. A summary of the results is shown in Table 2.

For Scenario A, GAM-PELT benchmarks well against the marginal approach with both methods correctly estimating zero changepoints in 98% of replicates. In Scenario B, where each spatial location has a different AR component, GAM-PELT performs slightly worse; estimating 1–4 false changepoints in 8% of replicates compared to only 1.64% using the marginal approach. For Scenarios C–F, GAM-PELT correctly estimates zero changepoints in 100%

Table 3 Summary of scenarios where changepoints are introduced into the simulated data

Scenario	Time	Space
1a	<i>All locations change:</i> AR1, Mean and Variance	<i>No change:</i> Constant
1b	<i>All locations change:</i> AR1, Mean and Variance	<i>No change:</i> All random
1c	<i>All locations change:</i> AR1, Mean and Variance	<i>No change:</i> Correlated
2a	<i>1 location changes:</i> AR1, Mean and Variance	<i>No change:</i> Constant
2b	<i>1 location changes:</i> AR1, Mean and Variance	<i>No change:</i> All random
2c	<i>1 location changes:</i> AR1, Mean and Variance	<i>No change:</i> Correlated
3a	<i>No change:</i> AR1, Mean and Variance	<i>Change:</i> Constant → Constant
3b	<i>No change:</i> AR1, Mean and Variance	<i>Change:</i> All random → All random
3c	<i>No change:</i> AR1, Mean and Variance	<i>Change:</i> Correlated → Correlated
4a	<i>Change:</i> Random draw (Can include no change)	<i>Change:</i> Random draw (Can include no change)
4b	<i>Change:</i> Random draw	<i>Change:</i> Random draw

of replicates run for each scenario. This slightly outperforms the marginal approach which demonstrates some evidence of over-fitting with false changepoints estimated in a small number of replicates (0.62–0.96%).

3.2 Temporal changes

We first evaluate GAM-PELT in terms of the ability to detect changepoints where there is a change in the temporal structure of the dataset only (Scenarios 1a–c and 2a–c in Table 3). Full details of the parameter settings are given in the Supplementary Material. The results are shown in Fig. 1.

For Scenario 1, GAM-PELT outperforms the marginal approach in all scenarios, correctly identifying a greater proportion of the true changepoints alongside a lower proportion of false positives. In contrast for Scenario 2, the marginal approach is shown to outperform GAM-PELT. This is expected as the change is at a single spatial location and the GAM-PELT parameter estimates are unlikely to change significantly due to this. Conversely, the marginal approach treats each spatial location in isolation (ignoring spatial dependencies) and therefore is better at capturing changes that impact single locations as in Scenario 2.

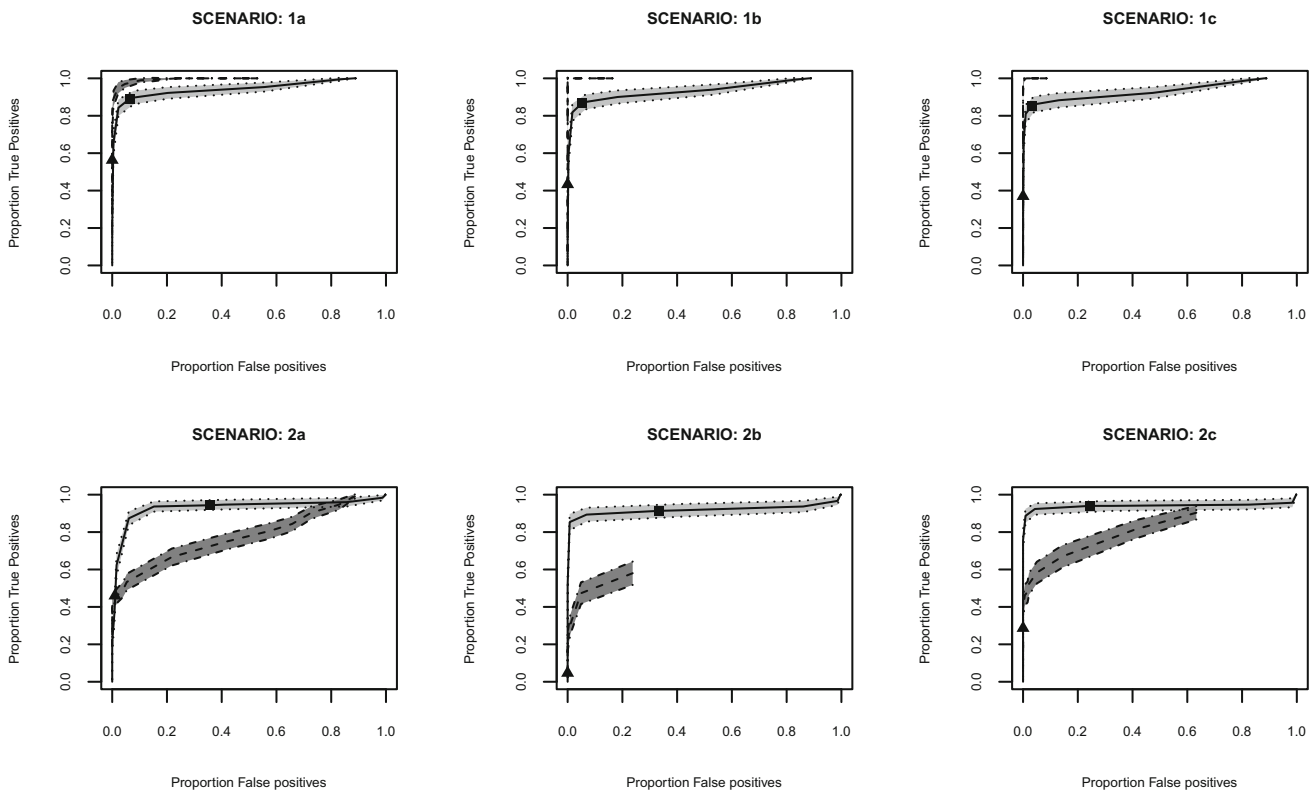


Fig. 1 Proportion of correctly identified changepoints against the proportion of falsely detected changepoints for Scenarios 1 (first row) and 2 (second row). GAM-PELT: thick dashed black line and dark grey shading, marginal approach: thick solid black line and light grey shading.

ing. Shading is a 95% confidence interval. The triangle and the square represent the BIC penalties for GAM-PELT and marginal approaches respectively

3.3 Spatial changes

We now evaluate the method in terms of the ability to detect changepoints where there is a change in the spatial structure of the dataset only (Scenario 3 in Table 3). Full details of the parameter settings are given in the Supplementary Material. Figure 2 shows a summary of the results.

For the detection of changes in the spatial structure, both methods performed well at detecting the timing of the changepoints, however, in all scenarios the GAM-PELT method was shown to outperform the marginal approach, once again showing lower proportions of false positives. The scenarios where GAM-PELT tends to perform much better are 3b (all random) and 3c (structured correlation). Recall that the marginal approach does not take account of the spatial structure.

3.4 Spatio-temporal changes

The final set of simulations evaluate the ability to detect changepoints where both the spatial and/or the temporal structure of the dataset changes between segments (Scenario 4 in Table 3). Full details of the parameter settings are given

in the Supplementary Material. Figure 3 shows a summary of the results.

For Scenario 4a (where no change is an option between changepoints) both methods show similar performance at detecting the timing of the changepoints. However, the GAM-PELT method is shown to perform slightly better; detecting a greater proportion of the true changepoints for fewer false positives. The stronger performance of GAM-PELT is highlighted for scenario 4b (where there is always a change of some type between changepoints) with this approach showing a higher proportion of true positives and a marked reduction in false positives over the marginal approach. This scenario is the most likely to be seen in practice.

3.5 Comparison to composite likelihood approach

Finally we compare the performance of GAM-PELT against the composite likelihood-minimum description length (CLMDL) approach proposed by Zhao et al. (2024) which also uses PELT for the multiple changepoint search. For this scenario we adopt the same four parameter autoregressive spatial model utilised in the simulation studies of Zhao et al. (2024) and simulate on an 8 by 8 regular two-dimensional


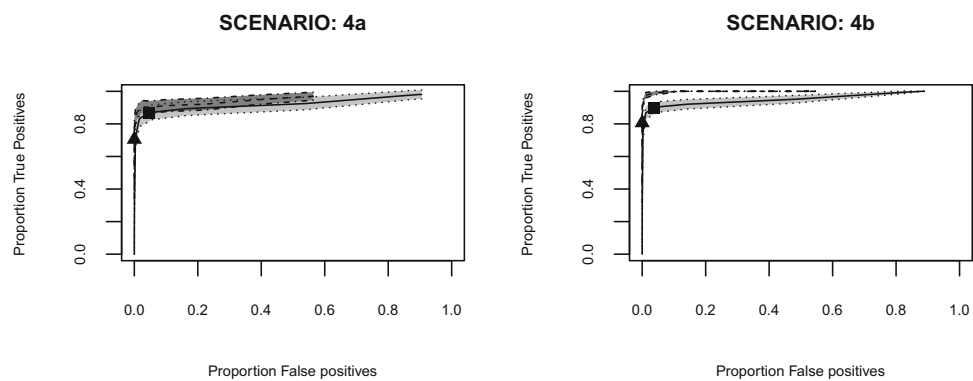


Fig. 2 Proportion of correctly identified changepoints against the proportion of falsely detected changepoints for Scenario 3. GAM-PELT: thick dashed black line and dark grey shading, marginal approach: Thick

solid black line and light grey shading. Shading is a 95% confidence interval. The triangle and the square represent the BIC penalties for GAM-PELT and marginal approaches respectively

Fig. 3 Proportion of correctly identified changepoints against the proportion of falsely detected changepoints for Scenario 4. GAM-PELT: thick dashed black line and dark grey shading, marginal approach: thick solid black line and light grey shading. Shading is a 95% confidence interval. The triangle and the square represent the BIC penalties for GAM-PELT and marginal approaches respectively



grid (with a grid spacing of 0.25 to simulate a real geographic grid) with 100 time points. We define a single true changepoint at $t = 50$, with a change in the signal strength of 0.3 in both the spatial and temporal components of the model after the changepoint. Both the GAM-PELT and CLMDL approaches are run for 100 replicates using their default settings and set to detect a minimum segment length of 20. Finally we run a no change scenario to evaluate both methods during situations of no change. The results are presented in Table 4.

The CLMDL approach is shown to capture the timing of the true changepoint in 90% of replicates, slightly outperforming GAM-PELT which captures the true changepoint in 80% of replicates. However, both methods demonstrate evidence of overfitting, with GAM-PELT less prone to this. For the no change scenario, GAM-PELT correctly estimates zero changepoints in 65% of replicates compared to only 6% using CLMDL. Here, CLMDL shows greater evidence of overfitting by estimating greater than 2 changepoints in 76% of replicates compared to 9% for GAM-PELT. We do however note here that we are reporting the performance of each method using their default penalties which take different approaches to using the penalty in the PELT algorithm (BIC vs MDL). Both approaches could benefit from employing

smaller penalties to reduce overfitting. Finally, GAM-PELT is shown to complete the 100 replicates for each scenario in around 0.3h which is approximately 23 times quicker than CLMDL which takes around 6.8h. This is majorly due to the difference in computation time for evaluating the likelihood in the two different models.

4 Data application

The GAM-PELT method was applied to air quality (AQ) station data from the United Kingdom (UK) Automatic, Urban and Rural Network (AURN). This network of 175 monitoring sites around the UK provides measurements of several key air pollutants at a frequency of up to 1h. More details about the data can be found in the supplementary material. The period 1st February - 31st August 2020 (213 days) was chosen as this covers the timeline of the UK's first nationwide COVID-19 lockdown, whereby impacts on pollutant concentrations would be expected to be seen in some effect at all monitoring locations. We focus on 2 primary (directly emitted) pollutants namely nitrogen dioxide (NO₂; measured at 74 spatial locations) and particulate matter of size smaller than 2.5 micron (PM_{2.5}; 30 spatial locations), and 1 sec-

Table 4 Comparison of GAM-PELT and CLMDL under different scenarios

Method	Scen	% Cpt Capture	% m = 0	% m = 1	% m >= 2	Run time
CLMDL	no cng	–	6	18	76	6.7 h
	1 cpt	90	0	11	89	6.8 h
GAM-PELT	no cng	–	65	26	9	0.34 h
	1 cpt	80	4	24	72	0.29 h

The true changepoint is classed as captured if the method estimates a changepoint within 10 of the true changepoint. The run time is based on running the 100 replicates in parallel using 7 cores on an intel i7 processor

ondary pollutant (formed in the atmosphere and thus behaves differently), namely ozone (O₃; 30 spatial locations). Here, the data was aggregated to daily averages which provided complete time series for all pollutants at the respective locations. If there were incomplete time series one could add an appropriate missing data handling procedure to the GAM fit within each segment. GAM-PELT was run with default settings (including BIC penalty) with the exception of the minimum segment length which was set to 15 days. A summary of the output is shown in Fig. 4.

GAM-PELT detects common changepoints at all spatial locations on the 26th March 2020 for O₃, 21st March 2020 for PM_{2.5}, and 27th March 2020 for NO₂, which correspond to the days around the nationwide UK lockdown on the 23rd March 2020. When the lockdown was introduced there was a sudden reduction in travel to work and other economic activity, and therefore an associated reduction in the emission of air pollutants that would be seen UK-wide. The changepoint for PM_{2.5} occurs slightly before the nationwide lockdown however, in the week before the national lockdown many people started working from home as a precaution which could account for an earlier change in particulate emissions. Figure 4 also shows the spatial components of the underlying GAM model before and after the onset of the lockdown period. Here, particularly for NO₂ and O₃, there is a noticeable shift in the nationwide spatial distribution of the pollutants when the lockdown commenced. Finally, changepoints that could be attributed to the first events of lifting the UK lockdown (Phased re-opening of schools from the 1st June 2020) are also detected for all pollutants (4th June 2020 for O₃, 11th June 2020 for PM_{2.5} and 7th June 2020 for NO₂).

5 Conclusion

We have developed a new spatio-temporal changepoint detection method (GAM-PELT) that can detect changes in spatially linked multivariate time series data. This method is implemented by utilising a generic GAM model (fitted on the spatial location and observed time of the data) in conjunction with the PELT search algorithm to detect changes

in the underlying spatio-temporal dependencies between the time series. When compared to a marginal approach (where a univariate model is applied to each spatial location in

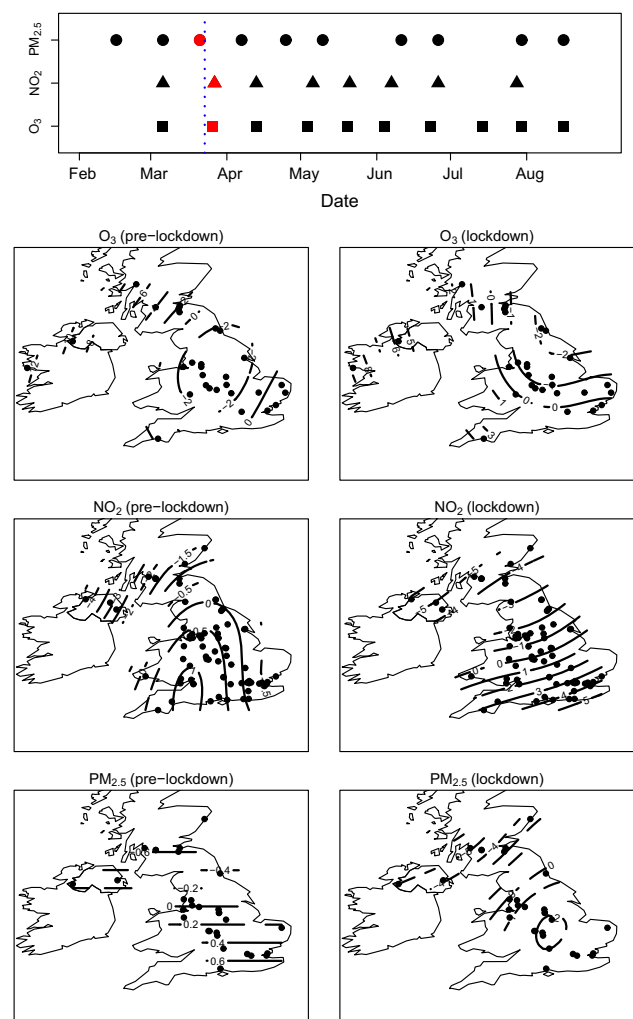


Fig. 4 Top plot: GAM-PELT changepoints across AURN stations measuring O₃ (squares), NO₂ (triangles) and PM_{2.5} (circles). The changepoints most likely associated with the start of the UK lockdown (23rd March 2020—dashed blue line) are shown in red. Maps: spatial component of GAM before (left column) and after (right column) the start of lockdown along with locations of AURN measurement stations for each pollutant. (Colour figure online)

isolation), the GAM–PELT method is shown to perform more efficiently in simulation studies at detecting the true changepoints and demonstrates less evidence of over-fitting. Furthermore, when treating each location in isolation, if 2 changepoints occur at different locations at the same time step, the multivariate power cannot be leveraged. As the GAM–PELT approach explicitly models the spatio-temporal dependencies between locations, our approach can detect common changepoints across the entire network of points. The effectiveness of the method was demonstrated through an application to an air quality dataset over the UK, where GAM-PELT was able to detect changepoints that may be linked to the onset and gradual lifting of the UK COVID-19 lockdown in 2020. Finally, when benchmarked against the existing state-of-the-art CLMDL approach of Zhao et al. (2024), GAM–PELT is shown to perform better at detecting the timing and number of true changepoints whilst demonstrating a runtime that is over 20 times faster.

It is important to note that any changepoint approach is sensitive to the model and penalty choices (C in (5)) made. Whilst not considered here, if the GAM model form is not constructed to be sensitive to the underlying changes within a given dataset, then the approach is unlikely to identify changepoints. In practice, slight overestimate of, for example, spline or tensor orders is preferable to underestimation, depending on which coefficients the change manifests within. Similarly for the penalty, C , if it is set too small then spurious changepoints may be detected, equally, too large and changepoints may be missed. Typical choices for C include the SIC/BIC (used here), MBIC (Zhang and Siegmund 2007), and data-driven methods based on the steepness of a scree-type plot (Lavielle 2005), or supervised learning (Hocking et al. 2013).

Supplementary information

This article has accompanying supplementary material dictating the full details used in the various simulation studies conducted in this manuscript.

Supplementary Information The online version contains supplementary material available at <https://doi.org/10.1007/s11222-024-10478-6>.

Acknowledgements MH acknowledges funding from the UK Engineering and Physical Sciences Research Council (EPSRC) under grant number EP/R01860X/1 and the UK Natural Environment Research Council (NERC) under grant number NE/R016429/1. RK acknowledges funding from the UK Engineering and Physical Sciences Research Council (EPSRC) under grant numbers EP/R01860X/1 and EP/T021020/1 and the UK Natural Environment Research Council (NERC) under Grant Number NE/T012307/1.

Author contributions MH designed and conceived the study, conducted the analysis and led the writing of the manuscript. RK designed and

conceived the study and commented and edited the first draft of the manuscript. All authors read and approved the final manuscript.

Code Availability The code to run the GAM-PELT method and reproduce the examples presented in this paper is available in the following github repository: <https://github.com/mjhollaway/GAM.PELT.git>

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

References

- Bakar, K.S., Sahu, S.K.: spTIMER: Spatio-temporal bayesian modeling using R. *J. Stat. Softw.* **63**(15):1–32 (2015). <https://doi.org/10.18637/jss.v063.i15>. <https://www.jstatsoft.org/index.php/jss/article/view/v063i15>
- Beaulieu, C., Killick, R.: Distinguishing trends and shifts from memory in climate data. *J. Clim.* **31**(23), 9519–9543 (2018). <https://doi.org/10.1175/JCLI-D-17-0863.1>. <https://journals.ametsoc.org/view/journals/clim/31/23/jcli-d-17-0863.1.xml>
- Beaulieu, C., Killick, R., Ireland, D., et al.: Considering long-memory when testing for changepoints in surface temperature: A classification approach based on the time-varying spectrum. *Environmetrics* **31**(1), e2568 (2020). <https://doi.org/10.1002/env.2568>. <https://onlinelibrary.wiley.com/doi/abs/10.1002/env.2568>, e2568 env-18-0130.R2
- Caron, F., Doucet, A., Gottardo, R.: On-line changepoint detection and parameter estimation with application to genomic data. *Stat. Comput.* **22**(2), 579–595 (2012). <https://doi.org/10.1007/s11222-011-9248-x>
- Creswell, R., Robinson, M., Gavaghan, D., et al.: A Bayesian nonparametric method for detecting rapid changes in disease transmission. *J. Theor. Biol.* (2023). <https://doi.org/10.1016/j.jtbi.2022.111351>
- Dette, H., Quanz, P.: Detecting relevant changes in the spatiotemporal mean function. *J. Time Ser. Anal.* **44**(5–6), 505–532 (2023). <https://doi.org/10.1111/jtsa.12674>
- Finley, A.O., Banerjee, S.: Bayesian spatially varying coefficient models in the spBayes R package. *Environ. Model. Softw.* **125**, 104608 (2020). <https://doi.org/10.1016/j.envsoft.2019.104608>. <https://www.sciencedirect.com/science/article/pii/S1364815219310412>
- Finley, A.O., Banerjee, S., Gelfand, A.E.: spBayes for large univariate and multivariate point-referenced spatio-temporal data models. *J. Stat. Softw.* **63**(13):1–28 (2015). <https://doi.org/10.18637/jss.v063.i13>. <https://www.jstatsoft.org/index.php/jss/article/view/v063i13>
- Gallagher, C., Lund, R., Robbins, M.: Changepoint detection in daily precipitation data. *Environmetrics* **23**(5), 407–419 (2012). <https://doi.org/10.1002/env.2146>
- Gromenko, O., Kokoszka, P., Reimherr, M.: Detection of change in the spatiotemporal mean function. *J. R. Stat. Soc. Ser. B (Stat. Methodol.)* **79**(1):29–50 (2017). <http://www.jstor.org/stable/44681761>

- Hahn, G., Fearnhead, P., Eckley, I.A.: Bayesproject: fast computation of a projection direction for multivariate changepoint detection. *Stat. Comput.* **30**(6), 1691–1705 (2020). <https://doi.org/10.1007/s11222-020-09966-2>
- Hocking, T., Rigaiill, G., Vert, J.P., et al.: Learning sparse penalties for change-point detection using max margin interval regression. In: Dasgupta, S., McAllester, D. (eds.) *Proceedings of the 30th International Conference on Machine Learning, Proceedings of Machine Learning Research*, vol 28. PMLR, Atlanta, Georgia, USA, pp. 172–180 (2013). <https://proceedings.mlr.press/v28/hocking13.html>
- Killick, R., Fearnhead, P., Eckley, I.A.: Optimal detection of change-points with a linear computational cost. *J. Am. Stat. Assoc.* **107**(500), 1590–1598 (2012). <https://doi.org/10.1080/01621459.2012.737745>
- Killick, R., Beaulieu, C., Taylor, S., et al.: EnvCpt: Detection of structural changes in climate and environment time series. (2021) <https://CRAN.R-project.org/package=EnvCpt>, R Package Version 1.1.3
- Laird, N.M., Ware, J.H.: Random-effects models for longitudinal data. *Biometrics* **38**(4):963–974 (1982). <http://www.jstor.org/stable/2529876>
- Lavielle, M.: Using penalized contrasts for the change-point problem. *Signal Process.* **85**(8), 1501–1510 (2005). <https://doi.org/10.1016/j.sigpro.2005.01.012>, <https://www.sciencedirect.com/science/article/pii/S0165168405000381>
- Liehrmann, A., Delannoy, E., Launay-Avon, A., et al.: DiffSegR: an RNA-seq data driven method for differential expression analysis using changepoint detection. *NAR Genom. Bioinf.* **5**(4):lqad098 (2023). <https://doi.org/10.1093/nargab/lqad098>. <https://academic.oup.com/nargab/article-pdf/5/4/lqad098/52946129/lqad098.pdf>
- Lindstrom, M.J., Bates, D.M.: Newton-raphson and em algorithms for linear mixed-effects models for repeated-measures data. *J. Am. Stat. Assoc.* **83**(404):1014–1022. (1988) <http://www.jstor.org/stable/2290128>
- Lowther, A.P., Killick, R., Eckley, I.A.: Detecting changes in mixed-sampling rate data sequences. *Environmetrics* **34**(1), e2762 (2023). <https://doi.org/10.1002/env.2762>. <https://onlinelibrary.wiley.com/doi/abs/10.1002/env.2762>
- Lund, R., Wang, X.L., Lu, Q.Q., et al.: Changepoint detection in periodic and autocorrelated time series. *J. Clim.* **20**(20), 5178–5190 (2007). <https://doi.org/10.1175/JCLI4291.1>. <https://journals.ametsoc.org/view/journals/clim/20/20/jcli4291.1.xml>
- Ma, T.F., Yau, C.Y.: A pairwise likelihood-based approach for changepoint detection in multivariate time series models. *Biometrika* **103**(2), 409–421 (2016). <https://doi.org/10.1093/biomet/asw002>
- Ryan, S., Killick, R.: Detecting changes in covariance via random matrix theory. *Technometrics* **65**(4), 480–491 (2023). <https://doi.org/10.1080/00401706.2023.2183261>
- Tapsoba, Jd.D., Wang, C.Y., Zangeneh, S., et al.: Methods for generalized change-point models: with applications to human immunodeficiency virus surveillance and diabetes data. *Stat. Med.* **39**(8), 1167–1182 (2020). <https://doi.org/10.1002/sim.8469>
- Wood, S.: *Generalized Additive Models: An Introduction with R*, 2nd edn. Chapman and Hall/CRC, Boca Raton (2017)
- Younes, L., Albert, M., Moghekar, A., et al.: Identifying change-points in biomarkers during the preclinical phase of alzheimer's disease. *Front. Aging Neurosci.* (2019). <https://doi.org/10.3389/fnagi.2019.00074>
- Zhang, N.R., Siegmund, D.O.: A modified Bayes information criterion with applications to the analysis of comparative genomic hybridization data. *Biometrics* **63**(1), 22–32 (2007). <https://doi.org/10.1111/j.1541-0420.2006.00662.x>, <https://onlinelibrary.wiley.com/doi/abs/10.1111/j.1541-0420.2006.00662.x>
- Zhao, Z., Ma, T., Ng, W., et al.: A composite likelihood-based approach for change-point detection in spatio-temporal processes. *J. Am. Stat. Assoc.* (2024). <https://doi.org/10.1080/01621459.2024.2302200>
- Zheng, C., Eckley, I., Fearnhead, P.: Consistency of a range of penalised cost approaches for detecting multiple changepoints. *Electron. J. Stat.* **16**(2), 4497–4546 (2022). <https://doi.org/10.1214/22-EJS2048>

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