Title:

Change Detection of Air Quality Time-series using the R Package AQEval

Authors:

Karl Ropkins[1*], Anthony Walker[2], Ian Philips[1], Christopher E. Rushton[1], Tony Clark[2] and James E. Tate[1].

[*] Corresponding Author Karl Ropkins, email: k.ropkins@its.leeds.ac.uk.

[1] Institute for Transport Studies, University of Leeds, Leeds, LS2 9JT, UK.

[2] Joint Air Quality Unit, Department for Transport & Department for Environment, Food and Rural Affairs, Marsham Street London, SW1P 4DF, UK.

Highlights:

- R package for investigation of discrete changes in air quality time-series.
- Break-point/segment detection and quantification.
- Signal isolation by deseasonalisation, deweathering and background subtraction.

Graphical Abstract:

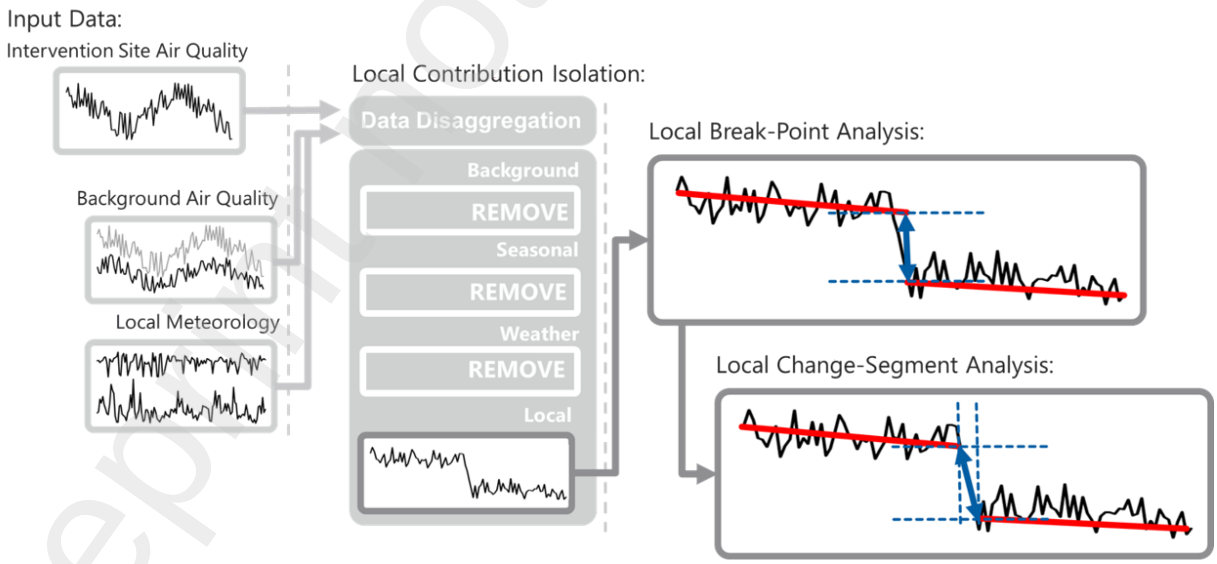

Abstract:

In this article, we introduce the R Air Quality Evaluation package AQEval. AQEval was developed for use by those tasked with the routine detection, characterisation and

quantification of discrete changes in air quality time-series, such as identifying the impacts of air quality policy interventions. The main functions use break-point/segment (BP/S) methods to first detect (as break-points) and then characterise and quantify (as segments) discrete changes in air quality time-series. Here, the objective is to introduce the main AQEval functions that provide robust and conservative estimates of change for new users, and a work-flow of methods and work-horse functions for those looking to fine-tune methods, e.g. through BP/S model optimisation and environmental signal isolation methods such as deseasonalisation, deweathering, and background subtraction.

Keywords:

R, AQEval, air quality time-series, change detection, break-point/segment.

1. Introduction:

Various change detection methods have been developed to investigate discrete changes in a wide range of time-series (see e.g. Reeves et al, 2007; Truong at al, 2020), and several R (R Core Team, 2020) packages have been produced for third-party use of related methods, e.g. bcp (Erdman & Emerson, 2007), changepoint (Killick et al, 2016), segmented (Muggeo, 2008), and strucchange (Zeileis et al, 2002).

So, why build a package specifically for air quality time-series change point analysis?

Authorities responsible for air quality management activities are typically required to implement and evaluate the air quality interventions they adopt (Public Health England, 2019). These interventions are often costly, disruptive, and even unpopular (Glazener & Khreis, 2019). Furthermore, robust confirmation of associated benefits is often hindered by the inherent variability of air quality data (Kelly et al, 2011), the influence of meteorology (Pearce et al, 2011; Grange & Carslaw, 2019), and the challenges of isolating one of several air pollution inputs at a monitoring site (Jones et al, 2012). Discrete change detection methods have previously been applied to the detection of large changes in air quality time-series Carslaw et al, 2006), and with, additional signal isolation, smaller changes associated with more modest interventions (Carslaw & Carslaw, 2007).

Those tasked with routine intervention assessment (e.g. local authority and government analysts, and those in transport planning consultancies who contract to government agencies) are typically highly skilled in a wide range of air quality monitoring and assessment activities, but are, realistically, unlikely to have the time and funding to develop in-house expertise in such specialist software. That said, many are already familiar with and using the R openair package (Carslaw & Ropkins, 2012). openair was developed specifically to work with the types of air quality data routinely collected by local authorities for assessment of regulatory air quality standards. It includes multiple functions for the routine analysis and visualisation of air quality data, including trend analysis using Thiel-Sen methods, but not the investigation and measurement of discrete changes in air quality time-series.

So, the key objectives in building AQEval were:

- To provide air quality professionals with an open package for routine air quality change point detection, characterisation and quantification, that facilitated access to such methods; And,
- To configure functions to handle air quality data in familiar formats using openair-like coding conventions to reduce the learning-curve associated with learning new software.

As part of the package development process, AQEval has already been used in two reported studies, Ropkins & Tate (2021) and Ropkins et al (2022), and this article provides an introduction to and discussion of associated package functions.

2. Package and Data Sources:

The AQEval package is freely available under General Public License (GPL) via conventional on-line R archives:

- The latest (stable) release version of AQEval is on the Comprehensive R Archive Network (CRAN) at https://CRAN.R-project.org/package=AQEval; And,
- The developers' version and code are publicly available on GitHub at (https://github.com/karlropkins/AQEval).

The release version of AQEval can be installed by download from R using the command install.packages("AQEval"), or as described in Appendix A if you are installing the developers' version or an AQEval (.zip or .tar.gz) package bundle.

Once installed, AQEval can be loaded conventionally within an R session:

```
R> require(AQEval) # or library(AQEval)
```

2.1. Data Sources:

All data sets used in this article are 1998 to 2005 1-hour resolution time-series of air quality and meteorological data, typical of that routinely collected and archived by air quality professionals as part of, e.g. regulatory air quality monitoring activities in the UK. These data sets are summarised as follows:

- **aq.my1** airborne nitrogen dioxide (NO_2), reported in $\mu g.m^{-3}$ from the London Marylebone Road monitoring station, and stored in the Automatic Urban and Rural Network (AURN) archive (AURN identifier code MY1). The MY1 site is located at one of busiest multi-lane inner city roadsides in the UK.
- **aq.ea2** NO_2 ($\mu g.m^{-3}$), wind speed ($m.s^{-1}$) and wind direction (degrees relative to North) from the Ealing Acton Town Hall station, and stored in the King's College London (KCL) network archive (KCL identifier code EA2). EA2 is also an urban roadside, but at a less heavily trafficked site by comparison to MY1.

- **aq.kc1** NO_2 ($\mu g.m_{-3}$) from the Kensington and Chelsea North Kensington Station, and stored in the KCL network archive (KCL identifier code KC1). KC1 is an urban background site, so more isolated from traffic emissions than the other two sites.

Although AQEval was developed with users of these and other similar archives in mind, there is no reason it cannot be used with any openair-friendly data. Typically, the data should be a single data frame (or similar object class), with a 'tidy' layout (as described in Wickham, 2014). So, each time-series should be included as a discrete column and each row should only contain simultaneous measurements. Although there are no restrictions to the types of data that can be included, one column should contain a conventional (POSIX class) time-series of date/time stamps and be named date so opeanir and AQEval functions can easily identify it. Likewise, wind speed and direction time-series should be named ws and wd, respectively.

See Ropkins and Carslaw (2012) for further discussion of openair-friendly data structures and Appendix B for example data sources.

2.2. Other Packages used in Examples:

As part of the demonstration of typical AQEval function usage, we also use several functions from other packages:

- **dplyr** we use dplyr functions and syntax for general data manipulation (Wickham et al, 2021). Although these are not the only options for similar pre-processing within R, arguably, code naming and structuring is highly intuitive and easier for those new to R to follow and apply (Broatch et al. 2019).
- **openair::timeAverage** we use openair function timeAverage to aggregate data sets by different time periods, e.g. to convert example 1-hour resolution data sets to 1, 3, 4 day or week (mean average) resolution data sets. Again, this is not the only option for similar pre-processing, but this is recommended here especially if data sets contain wind speed and direction data. See ?timeAverage in openair documentation for further details.

```
R> require(dplyr)
R> require(openair)
```

3. AQEval:

The AQEval Break-Point/Segment (BP/S) methods typically involve three steps: finding possible break-points, selecting the most likely break-points, and then quantifying these as either break-points or break-segments. The main package functions, quantBreakPoints and quantBreakSegments, are described in Section 3.1. to help practitioners getting started using the package and to introduce associated analytical strategies. Then, the workflow and options for those wanting to refine default

BP/S procedures are discussed in Section 3.2, and additional pre-processing options for those wanting to investigate smaller discrete changes in more complex time-series are discussed in Section 3.3.

3.1. Break-Points/Segments:

To find and quantify break-points, we employ functions based on the strucchange methods of Zeileis and colleagues (Zeileis et al. 2002, 2003, 2008) and independent change detection model testing as described in Section 3.2. As in most change detection strategies, strucchange adopts a rolling-window approach, assuming the first window (or data subset) is without change, building a statistical model of that, advancing the window, building a second model and comparing these, and so on, to identify most likely points of change in a larger data-series. The strucchange methods, however, test change using a statistical measure of difference between linear regression models of these windows, to provide a robust prediction of discrete change even when applied to time-series that are also subject to more gradual and persistent change (Figure 1). While no current change/break-point method completely prevents false-positives and false-negatives, this approach means strucchange is arguably better suited to use with air quality time-series, where changes rarely occur in isolation, than change detection methods based on assumptions of discrete static states before and after change events.

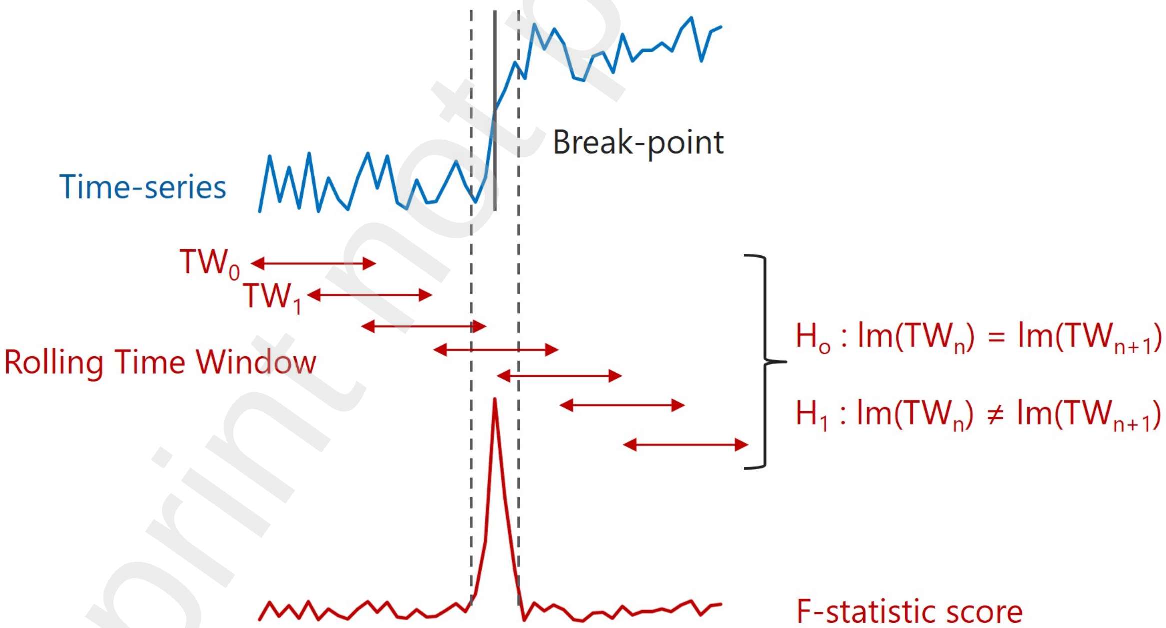

Figure 1: The strucchange break-point detection strategy. Here, TW_0 is the initial 'no-change' time-window, and this and all subsequent rolling time-windows are compared with the next, applying F-Statistics and the NULL hypothesis that the linear model of each is not significantly different to the next.

The quantBreakPoints formal argument structure is:

```
quantBreakPoints(data, pollutant, breaks, ylab = NULL, xlab = NULL, pt.col = c("lightgrey", "darkgrey"), line.col = "red", break.col = "blue", show = c("plot", "report"), ...)
```

Where data is the data source, typically an openair-friendly data.frame or similar, pollutant is the name of the data time-series to be analysed. All other arguments are optional and include: xlab and ylab, which reset x- and y-axis labels if alternative labelling is required; pt.col, line.col and break.col which set the plot point, trend line and break-point colours; And, show sets the shown elements of the function output, by default the analysis report and plot.

To break-point test and quantify the aq.my1 NO_2 time-series at one-day resolution, we first average the 1-hour resolution data, then use the quantBreakPoints call[1]:

```
R> aq.my1.day <- timeAverage(aq.my1, "day")

R> quantBreakPoints(aq.my1.day, "no2",
 +         ylab="AURN Marylebone Road no2 [ug/m3] \n Break-Points")
```

The default (break-point) plot output is shown in Figure 2 and the local report, printed to the R console, is:

```
Using 1 of 2 suggested breaks: 2
2003-02-05 (2003-01-14 to 2003-02-11)
82.66->112.4;29.69 (36%)
```

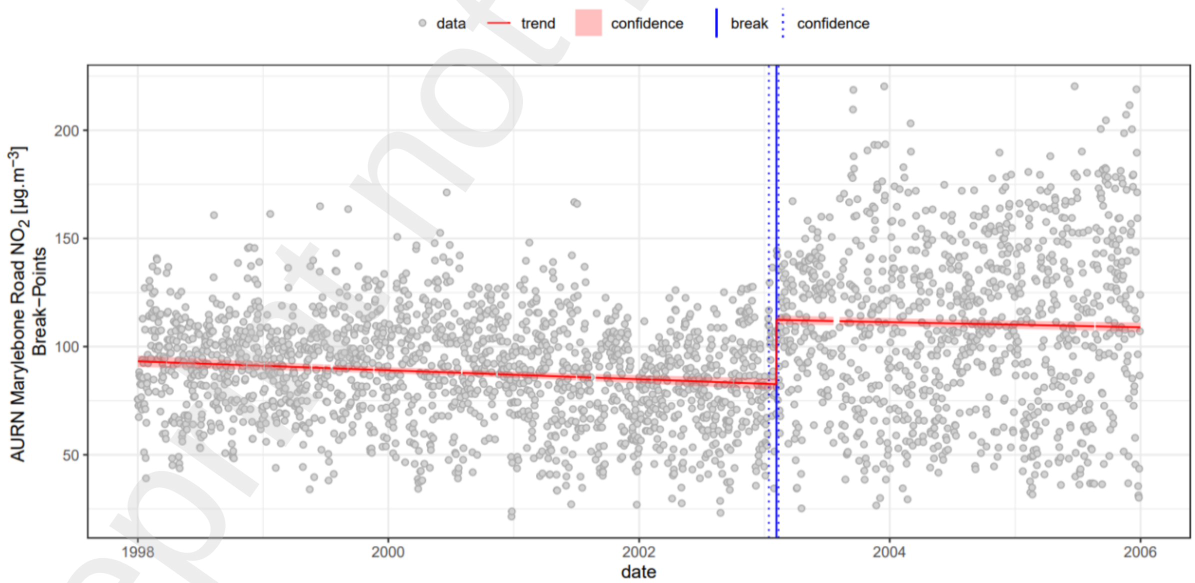

Figure 2: Marylebone Road 1998 to 2005 NO_2 time-series quantBreakPoints analysis using default setting.

[1] The ylab argument is optional, but is include as an example of routine plot annotation.

We are currently break-testing this data at one-day resolution because studies suggest that it provides a good compromise between processing speed and accuracy when working with several years of data (see e.g. Ropkins & Tate, 2021, and Ropkins et al, 2022). Most method evaluation to-date has been on time-series of similar lengths but this is not a specific requirement and other time-scales can also been studied. For example, in recent work we have also applied the methods to much shorter time-series, e.g. several months during the COVID-19 related lockdown, and early out-comes from this work suggest there may be advantages to working with data of 4-hour resolution and smaller time windows, when working with time-series of such lengths. Similarly, other analysis steps, e.g. signal isolation that may make use of hour-of-day measurements (see e.g. Section 3.3), may be better applied at the supplied 1-hour resolution.

This quantBreakPoints analysis shows two potential break-points were identified, but that a model applying only the second, occurring 2003-02-05 (95% confidence range 2003-01-14 to 2003-02-11), was considered most likely, and that it associated with an estimated discrete NO_2 increase of *ca* 30 μg.m^{-3} (from 82.66 to 112.4 μg.m^{-3}) or 36%.

Options to modify default break-point selections are discussed in Section 3.2.

Although this article is intended as an introduction to AQEval rather than an analysis of air quality in London between 1998 and 2005, it is perhaps worth noting that others, e.g. Carslaw and Carslaw (2007), have linked NO_2 increases at about these times with the introduction of EURO III vehicle technologies and associated London bus fleets upgrades.

Break-point detection methods assume an instantaneous switch from one set of conditions to another, and break-point quantification tends to be most reliable when the investigated changes tend towards this situation. However, few environmental changes are ever as rapid. The introduction of a cleaner fleet of vehicles, for example buses, often takes place through a series of phases, as newer, cleaner buses are delivered and take over services route-by-route. So, while break-point methods tend to be good at spotting where changes happen and providing an initial estimate of number and magnitude events, they are less reliable when it comes to more comprehensively characterising that change, especially if the change is more gradual. With this in mind, AQEval includes a second break quantification function, quantBreakSegments, which uses break-points identification as a starting-point but then tests the areas about these using the segmented methods of Muggeo (2003, 2008, 2017) to generate a rate-of-change model of the regions about the break-points (Figure 3).

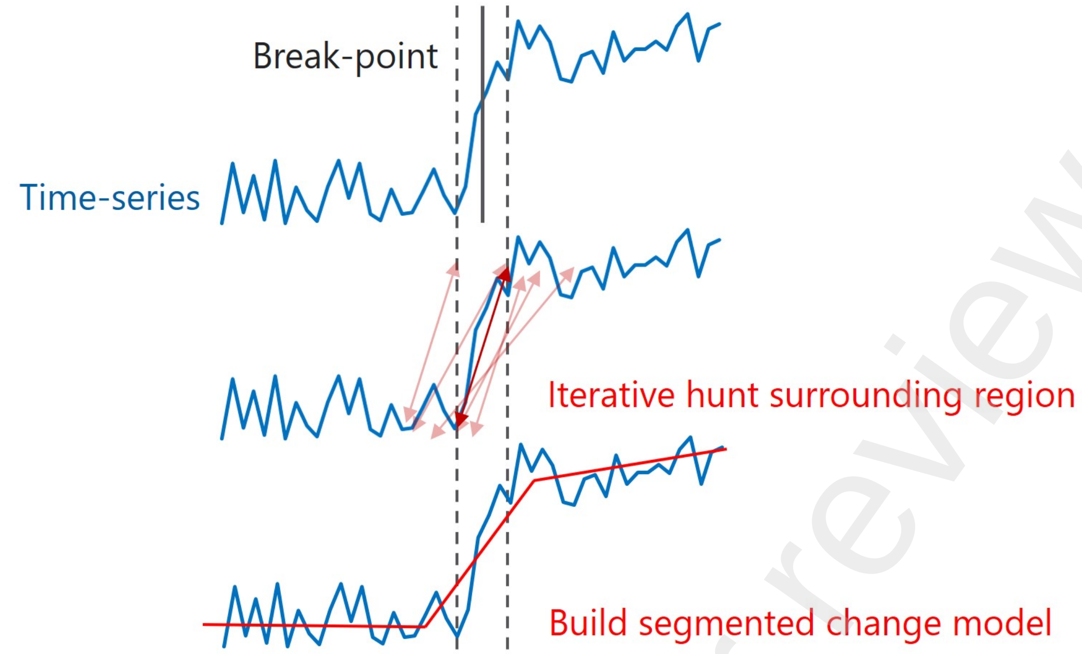

Figure 3: The AQEval Break segment fitting strategy. Here, the break-point (top black solid line) and confidence intervals (dotted lines) detected using strucchange methods are used as the initial segment ranges (red arrows), and segmented methods are then used to iteratively search about these (faded red lines) to build a break-segment model of associated change.

quantBreakSegments is applied in a similar fashion to quantBreakPoints, and tests potential breakpoints using the same methods but then generates a region of most likely change (or segment) model, assuming (2 × breaks) + 1) segments:

```
R> quantBreakSegments(aq.my1.day, "no2",
+     ylab="AURN Marylebone Road no2 [ug/m3] \n Break-Segments")
```

The break-segment plot output is shown in Figure 4 and the R console report is:

```
Using 1 of 2 suggested breaks: 2
1998-01-01 to 2003-01-11 (1836)
95.26->80.71;-14.55 (-15.28%)
2003-01-11 to 2003-02-19 (39)
80.71->105.9;25.17 (31.19%)
2003-02-19 to 2005-12-31 (1046)
105.9->115.3;9.371 (8.851%)
```

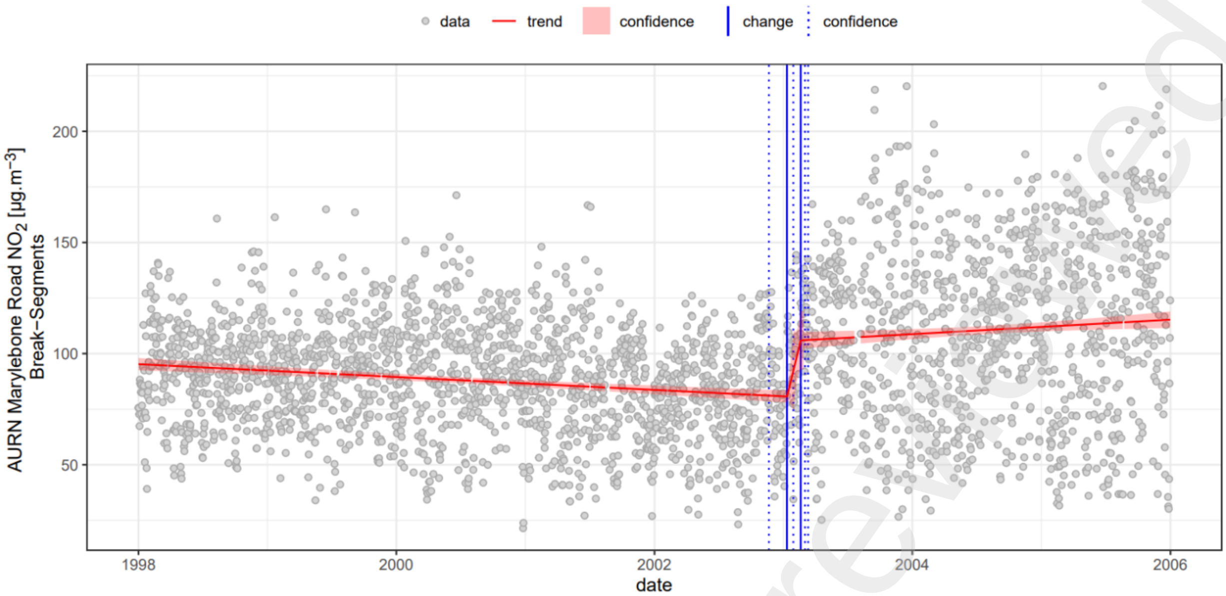

Figure 4: Marylebone Road 1998 to 2005 NO_2 time-series quantBreakSegments analysis using default setting.

Here, the break-segment model suggests a general decrease 1998 to 2002 (ca. 15% over 5 years), then a rapid step-change increase January to February 2003 (31% in about 1 month), followed by a more gradual increase March 2003 to start of 2006 (ca. 9% over 3 years). In this case, the break-segment associated with the main change event aligns very closely with the break-point confidence intervals (that does not always happen) and has a slightly lower magnitude (31%).

Because the segmented model allows for more freedom when fitting trends before and after the break-point event, it also suggests a more pronounced difference in trends before and after the region of change.

3.2. Workflow for Model Refinement

quantBreakPoints and quantBreakSegments apply a find > test > quantify analysis workflow: first finding potential break-points, then testing models using different combinations of these potential break-points, before quantifying a ‘best case’ using break-point and break-segment models, respectively. Users new to AQEval may prefer to leave this workflow to the code. However, users with knowledge of the data sets and/or local environment may prefer to review each of the steps and provide expert input. Those wishing to do so, can work with the full Break function sequence.

findBreakPoints uses the strucchange methods of Zeileis and colleagues (Zeileis et al. 2002, 2003) to identify potential change points within the supplied pollutant time-series. The formal argument structure is:

findBreakPoints(data, pollutant, h = 0.15, ...)

Where data and pollutant are as above, and addition argument h is the time-window size expressed as a proportion of the pollutant time-series length, so default 0.15 is equivalent to 15% of the aq.my1 NO_2 time-series in the workflow example:

```
R> bpts <- findBreakPoints(aq.my1.day, "no2")
```

The function includes missing data handling based on Zeileis (2017) and outputs a data.frame of identified break-points, assigned as bpts in the example, which the user can review, test or manually modify before passing on as argument breaks.

Independent break-point testing can be done using testBreakPoints, formal structure:

```
testBreakPoints(data, pollutant, breaks, ...)
```

For example:

```
testBreakPoints(aq.my1.day, "no2", bpts)
```

This expects data, pollutant and the output of findBreakPoints and identifies likely breakpoints using step-wise testing (Table 1):

Table 1: Marylebone Road 1998 to 2005 NO_2 time-series break-point test results (see main text for interpretation).

	Element(s)	breaks	significant	Adjusted R^2	suggest
1	2	1+2	FALSE	NA	
2	1	2	TRUE	0.1206	(<-)
3	1	1	TRUE	0.08388	
4	0	NA	FALSE	NA	

The initial (all break-points) model is accepted if it is statistically valid (or more specifically if all terms associated with individual break-points were all statistically significant, at $p<0.05$). If not, all model combinations discarding one of the initial break-points are built, tested and compared, and the statistically valid model with the highest correlation (all break-points individually statistically significant, $p<0.05$ and highest R) is accepted, or the process is repeated until a statistically valid model is obtained or all break-points were rejected. In this case, a one break-point model, using only the second of the proposed break-points, is recommended.

This step is included as an additional test of the potential break-points identified using strucchange Bayesian Information Criterion (BIC) testing methods, rather than an alternative. Additional tests were investigated in light of concerns raised about BIC by stucchange's authors (Zeileis et al. 2003), and the current additional method was selected on the basis of performance in simulation testing. That said, at this stage, this is presented as an empirical solution and arguably more work may yet be required on break-point selection.

The selected breaks can then be passed on to either of the quantification functions along with data and pollutant for either break-point or break-segment quantification:

```
R> #quantify break-points
R> ans <- quantifyBreakPoints(aq.my1.day, "no2", bpts)
R> #to use all proposed break-point or
R> #replace bpts with bpts[n:m,] or slice(bpts, n:m)
R> #to specify n to m break-points, etc
R> #for example bpts[2,] if using test recommendation

R> #quantify break-segments
R> ans <- quantBreakSegments(aq.my1.day, "no2", bpts)
R> #to use all proposed break-point or
R> #replace bpts with bpts[n:m,] or slice(bpts, n:m)
R> #to specify n to m break-points, etc
R> #for example bpts[2,] if using test recommendation
```

Both functions also allow some back passing of arguments. For example, if h is included in either quantification function call, it is passed back to findBreakPoints and overrides the default. Likewise, the extra argument test=FALSE can be used to pass all findBreakPoints detected break-points directly to the quantification function without local testBreakPoints testing. So, for example, the previous analysis could be simply modified to include all proposed break-points:

```
R> ans <- quantBreakSegments(aq.my1.day, "no2", test=FALSE,
+       ylab="AURN Marylebone Road no2 [ug/m3]
+       Break-Segments, using all proposed breaks")
```

Associated outputs are the break-segment plot shown as Figure 5 and the R console report:

```
Using all 2 suggested breaks
building 5 segments
1998-01-01 to 2000-07-14 (925)
90.49->93.53;3.041 (3.36%)
2000-07-14 to 2002-02-26 (592)
93.53->78.01;-15.52 (-16.59%)
```

2002-02-26 to 2003-01-03 (311)

78.01->82.65;4.635 (5.942%)

2003-01-03 to 2003-03-05 (61)

82.65->106.3;23.63 (28.6%)

2003-03-05 to 2005-12-31 (1032)

106.3->115.1;8.846 (8.323%)

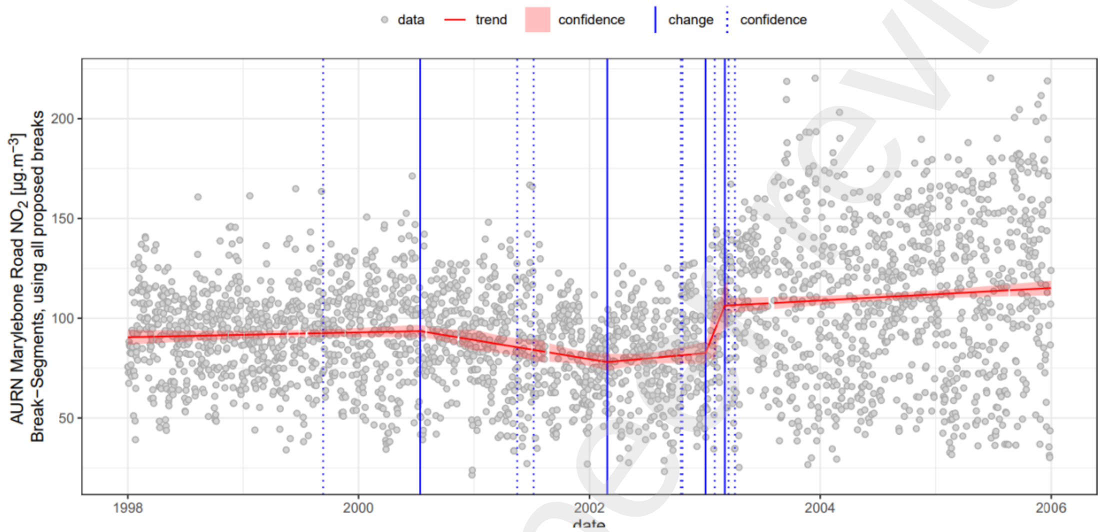

Figure 5: Marylebone Road 1998 to 2005 NO_2 time-series quantBreakSegments analysis using all proposed break-points (test=FALSE).

Here, both proposed break-points are used generating 5 segments, and the earlier segment, associated with the break-point rejected by testBreakPoints, suggests a smaller and shallower decrease (15.5 µg.m^{-3}; 16.6 %), over about a year and a half (2000-07-14 to 2002-02-26) ahead of the main change event, an increase in early 2003 (29% 2003-01-03 to 2003-03-05) (Figure 5) highly similar to that seen in the previous analysis (*cf* Figures 4 and 5). Although, this earlier change is considered much less likely, it is maybe worth noting that this was about the time EURO 3 passenger car regulations were introduced and was a decrease, while the main more definite event was an increase and arguably better aligns with later changes in larger vehicle fleet regulations. Here, we also catch the quantBreakSegment output as ans. Outputs from both quantification functions can be captured in this fashion if expert users would like to explore the analyses further. (Please see AQEval package and the project website documentation for further details).

3.3. Local Signal Isolation

The methods work well in the above example because the changes are relatively large and other contributions to measurement variance are relatively small. In other cases, e.g. at sites with smaller contributions from road traffic, other contributions, e.g. the variance associated with weather conditions and contributions from other sources may hinder the detection of break-points associated with local events. If we look at the NO_2 time-series from the nearby Ealing site for the same period, we see:

```
R> #convert Ealing 1-hour data to 1-day average
R> aq.ea2.day <- timeAverage(aq.ea2, "day")
R> #Ealing break-points, using default settings
R> quantBreakPoints(aq.ea2.day, "no2",
+       ylab="KCL Ealing Acton Town Hall no2 [ug/m3]
+       Break-Points")
```

Associated outputs are the break-point plot shown as Figure 6 and the Console report:

```
Using 3 of 4 suggested breaks: 1,2,4
1999-09-22 (1999-04-06 to 2000-04-18)
65.33->49.77;-15.56 (-24%)
2001-05-25 (2001-01-19 to 2001-09-30)
59.68->48.41;-11.27 (-19%)
2003-12-19 (2003-08-29 to 2004-05-11)
62.72->50.83;-11.89 (-19%)
```

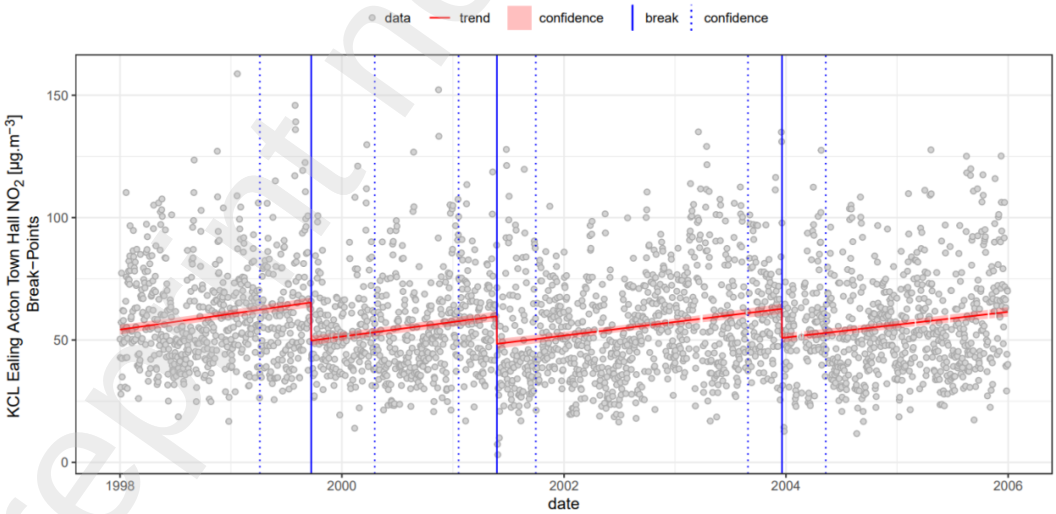

Figure 6: Ealing Acton Town Hall 1998 to 2005 NO_2 time-series quantBreakPoints analysis using default setting.

Here (Figure 6), NO_2 levels are lower and quantBreakPoints detects multiple possible breakpoints. However, these break-points are also almost regularly spaced out across the timeseries, generating a 'saw-tooth' pattern and very little net change across the full time-period, suggesting some degree of periodicity in the major source of change. Here it is perhaps worth noting that we would expect an environmental time-series to tend towards a regular series of break-points rather than exhibit exact regularly. For example, the change series here is not exactly yearly for several reasons:

- By default, we are applying a break-window of 15% (h=0.15), and that sets the window larger than one year on time-series of this size which may obscure/merge some changes happening at lower time resolutions.
- We are unlikely to ever see a perfect seasonal effect, i.e. an absolute association to day-of-year and exactly regular sequence of multi-year break-points, because the sources of seasonal variance are factors like the weather which, although closely associated with time-of-year, never themselves exhibit a perfect association with calendar date. So, often such events are easier to spot in ambient data using break-points rather than break-segments, which tend to emphasise the differences in scale of the events rather than their regularity-of-occurrence.

We can investigate this using different break-window sizes and time-series data ranges. For example, by increasing h so the time-window is much larger than one year, we can sometimes mask such seasonal (or near-seasonal) contributions if the change is towards the middle of our study time range (Figure 7):

```
R> #Ealing break-points, using longer time-window, h=0.3
R> quantBreakPoints(aq.ea2.day, "no2", h=0.3,
+ ylab="KCL Ealing Acton Town Hall no2 [ug/m3]
+ Break-Points (h=0.3)")
```

Associated outputs are the break-point plot shown as Figure 7 and the Console report:

```
Using 1 of 2 suggested breaks: 2
2002-09-25 (2002-06-02 to 2003-01-07) 49.05->62.13;13.09 (27%)
```

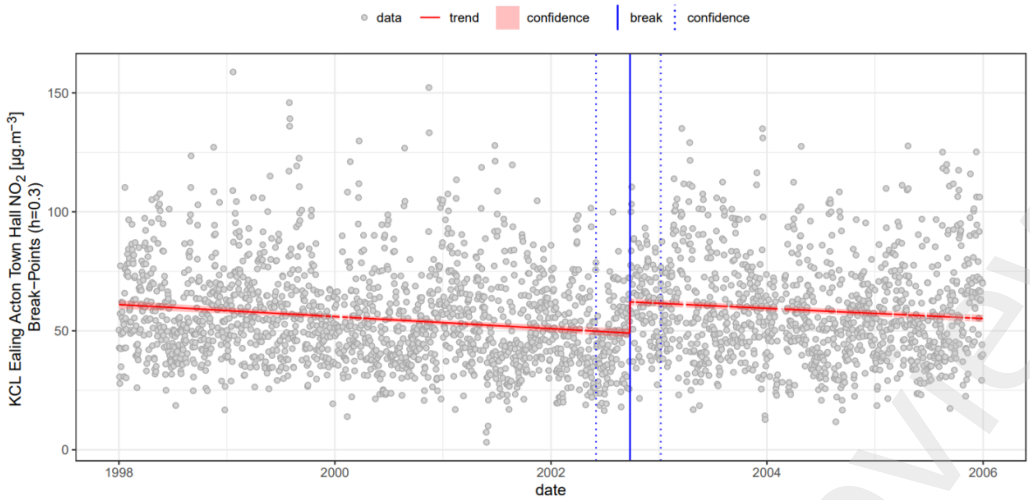

Figure 7: Ealing Acton Town Hall 1998 to 2005 NO_2 time-series quantBreakPoints analysis using larger time-window (h=0.3).

Similarly, by reducing h so the time window is less than one year and removing the break-point testing step, so we can see all proposed changes, we can enhance the detection of seasonal break-points (Figure 8)[2]:

Please note: This will take a while...

```
R> #Ealing break-points, using shorter time-window, h=0.05
R> #and all proposed breaks
R> quantBreakPoints(aq.ea2.day, "no2", h=0.05, test=FALSE,
+       ylab="KCL Ealing Acton Town Hall no2 [ug/m3]
+       Break-Points (h=0.05), using all proposed breaks")
```

Associated outputs are the break-point plot shown as Figure 8 and the Console report:

```
Using all 6 suggested breaks
1999-10-29 (1999-09-07 to 2000-01-15)
59.53->51.76;-7.775 (-13%)
2000-11-12 (2000-09-15 to 2000-12-21)
51.39->60.7;9.302 (18%)
2001-05-25 (2001-04-25 to 2001-07-01)
60.67->49.36;-11.31 (-19%)
2002-09-24 (2002-08-26 to 2002-10-19)
```

[2] Please note, If you are running the code yourself, this will take a while...

48.82->62.57;13.75 (28%)
2003-12-19 (2003-10-25 to 2004-03-22)
62.58->54.41;-8.174 (-13%)
2005-08-05 (2005-05-24 to 2005-09-11)
54.04->64.09;10.05 (19%)

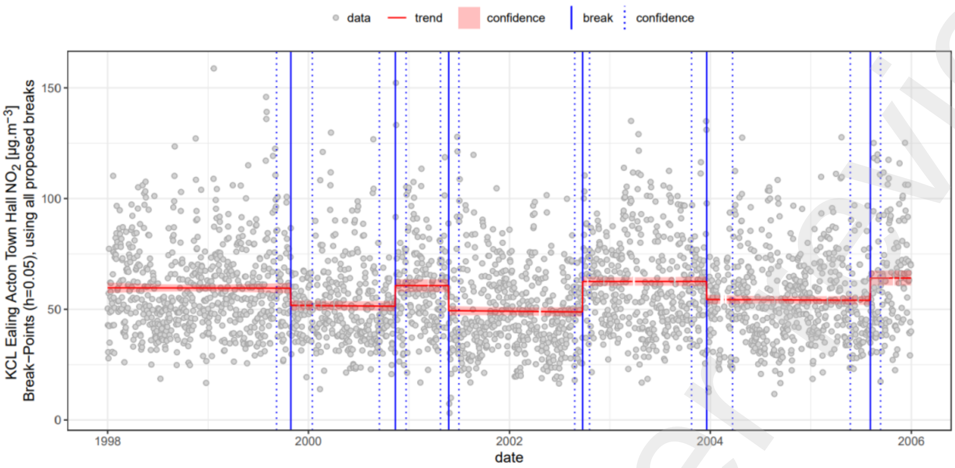

Figure 8: Ealing Acton Town Hall 1998 to 2005 NO_2 time-series quantBreakPoints analysis using smaller time-window (h=0.05) and all proposed break-points (test=FALSE).

Here, we see a pronounced cyclic up/down sequence to the proposed break-points, again not exactly regular but arguably consistent with annual meteorological cycles over a series of years in which, in terms of NO_2 air pollution, some were worse than others.

Various signal isolation strategies have been used to enhance components-of-interest within environmental time-series ahead of analyses:

- **Deseasonalisation** The removal of variance associated with regular frequency cycles within a time-series, e.g. hour-of-day, day-of-week, and week-of-year cycles (see e.g. Kendall and Stuart 1977). These are strictly proxies for environmental contributions varying at about these time-scales. The trends they map onto are not entirely regular but seasonal terms provide a practical alternative where more direct measures are not available and the analysed time-series are several times longer than the frequency cycles being subtracted. They can also be readily coded using date time-series of openair-friendly datasets.
- **Deweathering** The removal of variance associated with changes in meteorological conditions, such as wind speed and direction, air temperature and humidity (see e.g. Kuebler et al. 2001, Henneman et al. 2015). Sometimes referred to as 'weather normalisation' or 'meteorological detrending' in the

environmental research literature, these tend to be more direct and effective than seasonal terms where interfering variance is predominately weather driven, but requiring additional data, e.g. time-series of meteorological records.

- **Background Corrections** The removal of variance associated with (or direct subtraction of) levels measured at local background sites or predicted based on background modelling (see e.g. Stedman et al. 2006, Sayegh et al. 2016). For local processes, these can provide an enhanced measure of the foreground, often called the local 'increment' in the environmental research literature. The approach has repeatedly been proven highly effective for the quantification of roadside and urban increments, but is highly sensitive to the representativeness of the background data.
- **Filtering and Conditioning** The isolation of variance associated with indicators of the source-of-interest, e.g. source marker species or species ratios (see e.g. Cass 1998, Watson et al. 2008) or the removal of measurements when contributions from source-of-interest were expected to be minimal, e.g. up wind of a fixed source or out of operating hours of an industrial source (see e.g. Malby et al. 2013).

Here, we have divided these into four classes for the purposes of discussion. However, it is worth noting that these approaches are more often used in combination, with users selecting combinations based on the quality and representiveness of accessible local data.

To illustrate such environmental time-series signal isolation, we apply a deseasonalisation and deweathering (dSW) correction to the data using AQEval function isolateContribution and test the isolated contribution (what remains after we remove the dWS-associated variance from the time-series) for break-points and/or break-segments:

```
R> #dWS NO2 time-series at 1-hour resolution
R> aq.ea2$dws.no2 <- isolateContribution(aq.ea2, "no2")
R> #convert the data to 1-day resolution
R> aq.ea2.day <- timeAverage(aq.ea2, avg.time="day")
R> #break-point test
R> quantBreakPoints(aq.ea2.day, "dws.no2",
+        ylab="KCL Ealing Acton Town Hall no2 [ug/m3]
+        Break-Points, using dSW Signal Isolation")
```

Associated outputs are the break-point plot shown as Figure 9a and the Console report:

```
Using 1 of 2 suggested breaks: 2
2003-01-01 (2002-11-13 to 2003-02-11)
```

51.03->60.57;9.541 (19%)

(a) Break-Points

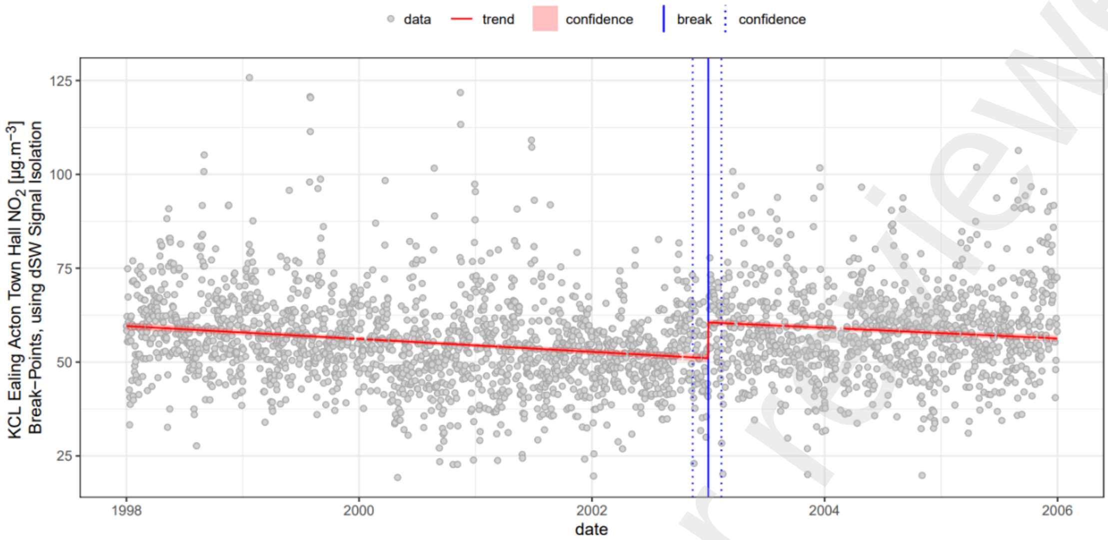

(b) Break-Segments

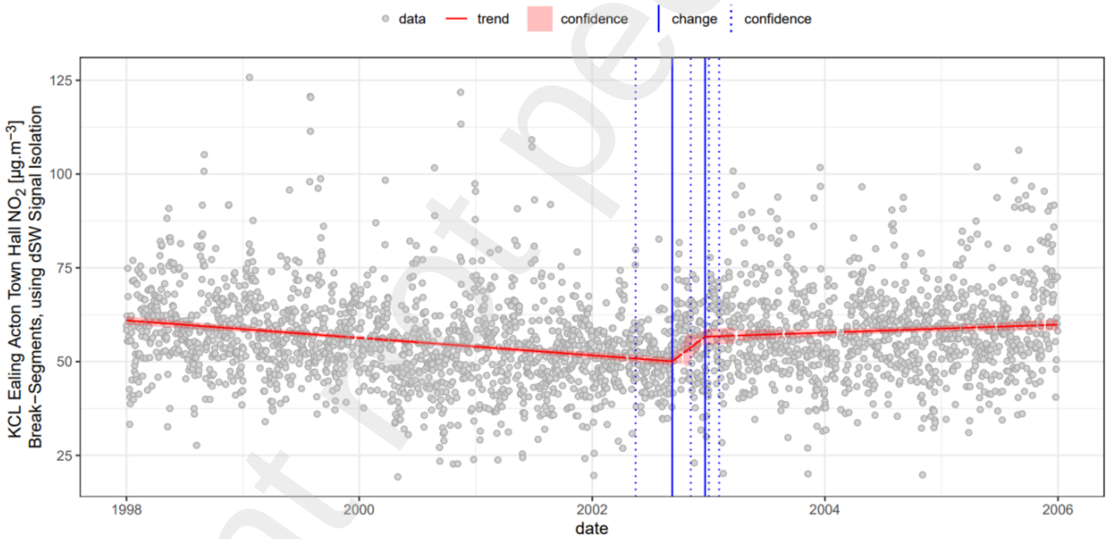

Figure 9: Ealing Acton Town Hall 1998 to 2005 NO2 time-series (a) quantBreakPoints and (b) quantBreakSegments analyses using simple (hour-of-day, Julian date, wind speed and wind direction) deseasonalisation and deweathering (dSW) signal isolation.

```
R> #break-segment test
R> quantBreakSegments(aq.ea2.day, "dws.no2",
+         ylab="KCL Ealing Acton Town Hall no2 [ug/m3]
+         Break-Segments, using dSW Signal Isolation")
```

Associated outputs are the break-segment plot shown as Figure 9b and the Console report:

```
Using 1 of 2 suggested breaks: 2
building 3 segments
1998-01-01 to 2002-09-09 (1712)
60.97->50.08;-10.9 (-17.87%)
2002-09-09 to 2002-12-21 (103)
50.08->56.65;6.57 (13.12%)
2002-12-21 to 2005-12-31 (1106)
56.65->59.84;3.195 (5.639%)
```

isolateContribution is a general-purpose 'local' contribution isolation function.

If isolateContribution is only supplied data and pollutant (and data contains date, pollutant, ws and wd time-series), by default it applies a simple dSW model in the form:

$$[poll] = s_1(j.day) + s_2(day.hour) + te_1(ws, wd) \qquad (1)$$

$$[poll]_{dsw} = ([poll] - [poll]_{model}) + \text{mean}([poll]) \qquad (2)$$

Where [*poll*] is the time-series of measurements of interest, typically concentrations of the pollutant requiring dSW, j.day and day.hour are the seasonal terms, Julian date (or day-of-year) and hour-of-day, respectively, and ws and wd are the weather terms, wind speed and wind direction, respectively. The contribution model (Equation 1) is fitted as a Generalised Additive Model (GAM) using the mgcv package (Wood 2017, Wood 2021) and thin-film splines for single input functions, s_1 and s_2 for day.hour and j.day, respectively, and a cubic-regression tensor, te_1 for the two-input ws and wd term. Finally, the mean-centred residual of the model is reported as the dSW output, $[poll]_{dsw}$ (Equation 2).

Here (Figure 9), we again see that single break-point/segment at the end of 2002/start of 2003, at about the same time as the main break-point in the first example data set. This indicates that the (hour-of-day, day-of-year, wind speed and wind direction) dSW corrected NO_2 change of the order of 9.5 and 6.6 $\mu g.m^{-3}$ (or 19% and 13%) using break-point and break-segment quantification, respectively.

Overly aggressive modelling will most likely remove or distort variance associated with the change of interest. So, we start with a relatively simple isolateContribution model. Should you need or want to, options to develop your model using include:

- **Using alternative weather data.** In the above example, we use archived meteorological measurements included in the KCL air quality data archive. Such data may not be available, may be of poor quality or sparsely available, you may have access alternative local data that you feel is more representative or wish to use an alternative source. Here, one alternative we have used in previous studies is the National Oceanic and Atmospheric Administration (NOAA) Integrated Surface Database (ISD), which can be downloaded using e.g. the R packages rnoaa (Chamberlain 2021) or worldmet (Carslaw 2021).

- **Adding additional weather-related parameters.** We would recommend using wind speed and direction if available and you are applying deweathering because pollution levels typically associate strongly with these two parameters. However, other work has indicated that additional meteorological measures, such as air temperature and rainfall can be helpful additions when deweathering. The isolateContribution argument deweather can be used to remove or to change the deweathering terms to the signal isolation.
- **Changing seasonal parameters.** Likewise, the deseasonalisation terms can be removed or changed using the isolateContribution deseason argument. However, here, it is important to note one distinction between deweathering and desesonalisation: deweathering is removing variance associated with meteorological change; deseasonalisation is removing variance of a given frequency. So, while deseasonalisation will remove some weather–related variance, especially if more direct measures are not available, it may also remove some portion of contributions from other sources with near-regular frequency patterns.
- **Using a background correction (BC).** If you have access to pollutant data from a suitable background site, you can calculate a conventional BC by subtracting this time-series from your study-site pollutant time-series to produce a local increment or apply it as a GAM term, .e.g. using the background argument in isolateContribution. This can be carried out in combination with dSW, or as a standalone BC by turning off the deseason and deweather options.
- **Accounting for other known contributions.** Where other sources are known and believed to be significant, it can sometimes be possible to reduce their influence if activity and/or emissions records are available as similar resolution time-series. These can, e.g. be added to the signal isolation step using the isolationContribution add.term argument.
- **Apply an alternative model** Users that are familiar with GAMs and mgcv syntax can use the isolateContribution argument formula to apply alternative signal isolation models. This overrides all other settings, so expects the alternative model formulae in full.

You can also use signal isolation of your own or produced using other packages.

However, when doing any environmental signal isolation, our current recommendation is to start with the pollutant of interest and no signal isolation and then add signal isolation as required, e.g. to remove variance from other sources potentially obscuring the change of interest, or to apportion contributions in cases where the magnitude a discrete change is expected to be influenced by several factors. In all cases it is most likely better to first test any additional inputs to ensure that they are not the potential source of any new changes observed when they are included in the signal isolation step.

As a simple example of the modification of the signal isolation step, we expanded the previous dSW to a dSW and BC using NO_2 from the nearby AURN Kensington Urban Background site (KC1) as our background:

```
R> #add background site no2 to data
R> select(aq.kc1, date, no2) %>%
+        rename(bg.no2=no2) %>%
+        right_join(aq.ea2, by="date") -> aq.ea2.bg
R> #dsw and background subtraction
R> aq.ea2.bg$dwsb.no2 <- isolateContribution(aq.ea2.bg,
+        "no2", background="bg.no2")
R> #convert 1-hour to 1-day
R> aq.ea2.bg.day <- timeAverage(aq.ea2.bg, avg.time="day")
R> #break-point test
R> quantBreakPoints(aq.ea2.bg.day, "dwsb.no2",
+        ylab="KCL Ealing Acton Town Hall no2 [ug/m3]
+        Break-Points, using dSW and Urban BC Signal Isolation")
```

Associated outputs are the break-point plot shown as Figure 10a and the Console report:

```
Using 1 of 2 suggested breaks: 2
2002-09-13 (2002-08-24 to 2002-09-20)
52.78->59.91;7.132 (14%)

R> #break-segment test
R> quantBreakSegments(aq.ea2.bg.day, "dwsb.no2",
+        ylab="KCL Ealing Acton Town Hall no2 [ug/m3]
+        Break-Segments, using dSW and Urban BC Signal Isolation")
```

Associated outputs are the break-segment plot shown as Figure 10b and the Console report:

```
Using 1 of 2 suggested breaks: 2
building 3 segments
```

1998-01-01 to 2002-09-14 (1717)

57.93->53.14;-4.791 (-8.27%)

2002-09-14 to 2002-09-15 (1)

53.14->56.32;3.181 (5.987%)

2002-09-15 to 2005-12-31 (1203)

56.32->59.93;3.614 (6.416%)

(a) Break-Points

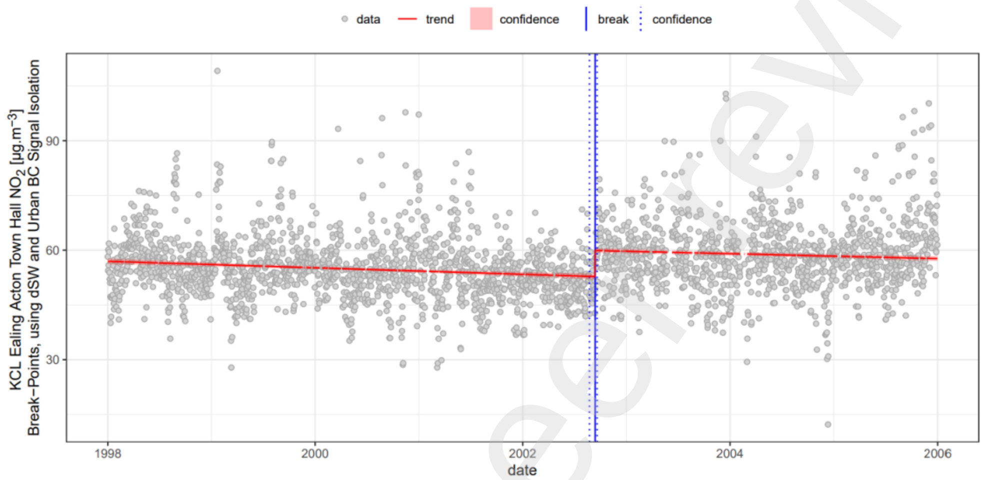

(b) Break-Segments

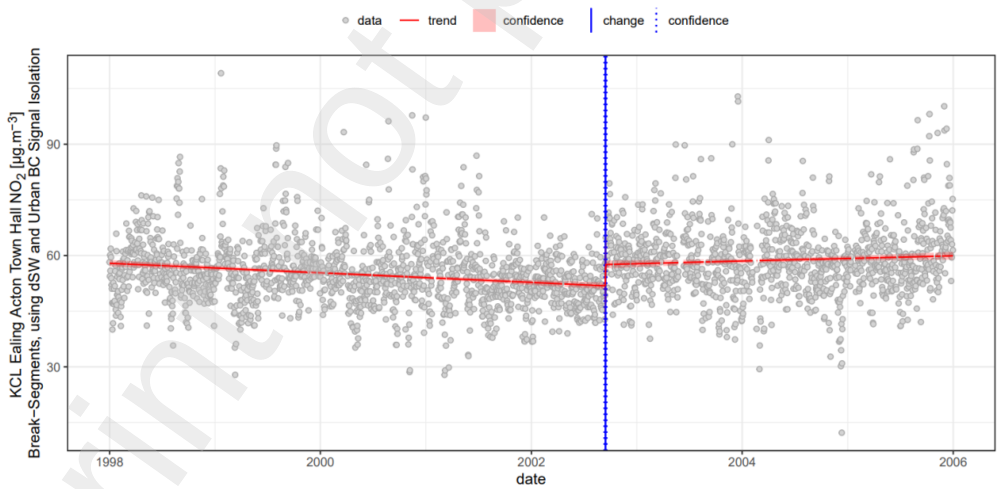

Figure 10: Ealing Acton Town Hall 1998 to 2005 NO2 time-series (a) quantBreakPoints and (b) quantBreakSegments analyses using simple (hour-of-day, Julian date, wind speed and wind direction) deseasonalisation and deweathering (dSW) signal isolation and background correction (BC) using NO2 data from the Kensington and Chelsea Urban Background site.

This expands Equation 1 to:

$$[poll] = s_1(j.day) + s_2(day.hour) + te_1(ws, wd) + s_3(background) \quad (3)$$

Where, s_3 is an additional spline term which is used to expand the previous dSW model to a dSW and BC model using the supplied background time-series.

This analysis indicates a dSW and BC NO_2 change event, this time more definitely in late 2002, and a local change (roadside by comparison to urban background) of about 7.1 μg.m^{-3} or 14% using break-points and 3.2 μg.m^{-3} or 6% using break-segments.

4. Summary and Discussion

Here, we have introduced the R AQEval package, provided background on the methodology, and demonstrated and discussed the use of its main functions.

We hope this article provides a helpful reference for those wishing to use AQEval to detect, characterise, and measure discrete changes in air quality time-series. For those interested in finding out more, further discussion and examples of these methods are provided in the AQEval package documentation and in associated online documentation (https://github.com/karlropkins/AQEval).

However, even this brief analysis provides some interesting insights:

- Comparing the dSW break-point results (Figure 9a and associated reports) with those from the previous longer time-window analysis (no dSW but h=0.3, Figure 8 and associated reports), we see a NO_2 difference of ca. 3.6 μg.m^{-3} (13.1-9.5 μg.m^{-3}). This suggests that about 27% (3.6/13.1 x 100) of NO_2 at this site may associate with the applied dSW terms. That said, ambient NO_2 chemistry is complex: A range of atmospheric processes drive its secondary production and degradation, and factors like wind speed and direction can influence pollutant accumulation in and/or dispersion from a monitoring site. So, it is important not to interpret this as a classic source, and assign this purely to, e.g., the weather. Instead it is perhaps more helpful to consider the weather a source/sink contribution, either increasing or reducing monitor site pollution levels depending on the prevailing conditions.
- Similarly, when comparing changes observed with and without BC (*cf* Figures 9 and 10 and associated reports), we see smaller reductions once the backgrounds are accounted for: 7.1 μg.m^{-3} by comparison to 9.5 μg.m^{-3} and 3.2 μg.m^{-3} by comparison to 6.6 μg.m^{-3}, for break-point and break-segment analyses, respectively. Again, we need to be cautious in our interpretation of this result because it could either reflect larger levels of secondary NO_2 at the background site, or it could indicate an poorly selected background, with a partial contribution from your change event. However, you could go on with AQEval to investigate changes at other sites and for other pollutants, and this just may help you to investigate these questions.

The package and associated research continue to be works in progress. We therefore recommend that users check package and on-line documentation for news and updates regarding changes to the package (see also Section C of the Supporting information). We would also welcome discussion about future options for the methods and third-party insights on the performance and sensitivity of the methods.

Likewise, we would happy to discuss collaborative projects applying the methods to both novel data sets and also well-understood historical time-series, because validation does not come from the reporting of a novel result, it comes from the confirmation of an independently established outcome.

Finally, although AQEval was developed as an environmental time-series analysis tool, other applications are also of interest. For example, with the signal isolation functions disabled, the BP/S analysis methods have already been applied to traffic counter records as part of a comparison of discrete changes in NO_2 and traffic flows during the first COVID-19 related lockdown (see e.g. Ropkins and Tate 2021), and this is just one area of interest we would like to investigate further in future.

Acknowledgements

This work was funded by the Department for Environment, Food and Rural Affairs (Defra). The authors also gratefully acknowledge contributions and input from colleagues at University of Leeds, Defra and IPSOS Mori, and input from the Defra/Department for Transport Joint Air Quality Unit (JAQU) Technical Independent Review Panel (T-IRP) as part of internal review, and from the editors as part of the external peer-review process at Environmental Software and Modelling.

The authors also gratefully acknowledge the work of the R core team and their many collaborators in developing and maintaining the open-source statistical language R and associated packages (http://www.r-project.org/).

The views and opinions expressed herein by the authors are their own and do not necessarily reflect those of UK Government or any agency thereof.

References

Broatch, J.E., Dietrich, S. and Goelman, D., 2019. Introducing data science techniques by connecting database concepts and dplyr, Journal of Statistics Education, 27(3), pp. 147-153, https://doi.org/10.1080/10691898.2019.1647768.

Carslaw, D.C., 2021. worldmet: Import Surface Meteorological Data from NOAA Integrated Surface Database (ISD). R package version 0.9.5, URL https://CRAN.R-project.org/package=worldmet.

Carslaw, D.C. and Carslaw, N., 2007. Detecting and characterising small changes in urban nitrogen dioxide concentrations. Atmospheric Environment, 41(22), pp.4723-4733. https://doi.org/10.1016/j.atmosenv.2007.03.034.

Carslaw, D.C. and Ropkins, K., 2012. openair - an R package for air quality data analysis. Environmental Modelling & Software, 27, pp.52-61. https://doi.org/10.1016/j.envsoft.2011.09.008.

Carslaw, D.C., Ropkins, K. and Bell, M.C., 2006. Change-point detection of gaseous and particulate traffic-related pollutants at a roadside location. Environmental science & technology, 40(22), pp.6912-6918. https://doi.org/10.1021/es060543u.

Cass, G.R., 1998. Organic molecular tracers for particulate air pollution sources. TrAC Trends in Analytical Chemistry, 17(6), pp.356-366. https://doi.org/10.1016/S0165-9936(98)00040-5.

Chamberlain, S., 2021. rnoaa: 'NOAA' Weather Data from R. R package version 1.3.8, https://CRAN.R-project.org/package=rnoaa.

Erdman, C. and Emerson, J. W., 2007. An R Package for Performing a Bayesian Analysis of Change Point Problems, Journal of Statistical Software, 23 (3), pp.1-13. http://www.jstatsoft.org/v23/i03/.

Glazener, A. and Khreis, H., 2019. Transforming our cities: best practices towards clean air and active transportation, Current environmental health reports, 6(1), pp.22-37, Springer, https://doi.org/10.1007/s40572-019-0228-1.

Grange, S.K. and Carslaw, D.C., 2019. Using meteorological normalisation to detect interventions in air quality time series. Science of the Total Environment, 653, pp.578-588. https://doi.org/10.1016/j.scitotenv.2018.10.344.

Henneman, L.R., Holmes, H.A., Mulholland, J.A. and Russell, A.G., 2015. Meteorological detrending of primary and secondary pollutant concentrations: Method application and evaluation using long-term (2000–2012) data in Atlanta. Atmospheric Environment, 119, pp.201-210. https://doi.org/10.1016/j.atmosenv.2015.08.007.

Jones, A.M., Harrison, R.M., Barratt, B. and Fuller, G., 2012. A large reduction in airborne particle number concentrations at the time of the introduction of "sulphur free" diesel and the London Low Emission Zone. Atmospheric Environment, 50, pp.129-138. https://doi.org/10.1016/j.atmosenv.2011.12.050.

Kelly, F., Anderson, H.R., Armstrong, B., Atkinson, R., Barratt, B., Beevers, S., Derwent, D., Green, D., Mudway, I. and Wilkinson, P., 2011. The impact of the congestion charging scheme on air quality in London. Part 1. Emissions modeling and analysis of air pollution measurements. Research Report (Health Effects Institute), (155), pp.5-71.

Kendall, M. and Stuart, A., 1977. "The advanced theory of statistics. Vol. 1: Distribution theory." London: Griffin.

Killick, R., Haynes, K., and Eckley, I.A., 2016. changepoint: An R package for changepoint analysis, R package version 2.2.2, https://CRAN.R-project.org/package=changepoint.

Kuebler, J., van den Bergh, H. and Russell, A.G., 2001. Long-term trends of primary and secondary pollutant concentrations in Switzerland and their response to emission controls and economic changes. Atmospheric Environment, 35(8), pp.1351-1363. https://doi.org/10.1016/S1352-2310(00)00401-5.

Malby, A.R., Whyatt, J.D. and Timmis, R.J., 2013. Conditional extraction of air-pollutant source signals from air-quality monitoring. Atmospheric Environment, 74, pp.112-122. https://doi.org/10.1016/j.atmosenv.2013.03.028.

Muggeo, V.M.R., 2003. Estimating regression models with unknown break-points, Statistics in Medicine, 22, pp.3055-3071.

Muggeo, V.M.R., 2008. segmented: an R Package to Fit Regression Models with Broken-Line Relationships, R News, 8(1), pp.20-25, https://cran.r-project.org/doc/Rnews/.

Muggeo, V.M.R., 2017. Interval estimation for the breakpoint in segmented regression: a smoothed score-based approach, Australian & New Zealand Journal of Statistics, 59, pp.311-322, https://doi.org/10.1111/anzs.12200.

Pearce, J.L., Beringer, J., Nicholls, N., Hyndman, R.J. and Tapper, N.J., 2011. Quantifying the influence of local meteorology on air quality using generalized additive models. Atmospheric Environment, 45(6), pp.1328-1336. https://doi.org/10.1016/j.atmosenv.2010.11.051.

Public Health England, 2019. Review of interventions to improve outdoor air quality and public health. https://assets.publishing.service.gov.uk/government/uploads/system/uploads/attachment_data/file/937341/Principal_interventions_for_local_authorities-air_quality_public_health.pdf.

R Core Team, 2020. R: A language and environment for statistical computing. R Foundation for Statistical Computing, Vienna, Austria. https://www.R-project.org/.

Reeves, J., Chen, J., Wang, X.L., Lund, R. and Lu, Q.Q., 2007. A review and comparison of changepoint detection techniques for climate data. Journal of applied meteorology and climatology, 46(6), pp.900-915. https://doi.org/10.1175/JAM2493.1

Ropkins, K. and Carslaw, D.C., 2012. openair - Data Analysis Tools for the Air Quality Community, The R Journal, 4(1), pp.20-29, https://doi.org/10.32614/RJ-2012-003.

Ropkins, K. and Tate, J.E., 2021. Early observations on the impact of the COVID-19 lockdown on air quality trends across the UK. Science of the Total Environment, 754, p.142374. https://doi.org/10.1016/j.scitotenv.2020.142374.

Ropkins, K., Tate, J.E., Walker, A. and Clark, T., 2022. Measuring the impact of air quality related interventions. Environmental Science: Atmospheres, 2(3), pp.500-516. https://doi.org/10.1039/D1EA00073J.

Sayegh, A., Tate, J.E. and Ropkins, K., 2016. Understanding how roadside concentrations of NOx are influenced by the background levels, traffic density, and meteorological conditions using Boosted Regression Trees. Atmospheric Environment, 127, pp.163-175. https://doi.org/10.1016/j.atmosenv.2015.12.024.

Stedman, J., Abbott, J., Willis, P. and Bower, J., 2006. Review of Background Air Quality Data and Methods to Combine these with Process Contributions. Environment Agency for England and Wales: Bristol.

Truong, C., Oudre, L. and Vayatis, N., 2020. Selective review of offline change point detection methods. Signal Processing, 167, p.107299. https://doi.org/10.1016/j.sigpro.2019.107299.

Watson, J.G., Antony Chen, L.W., Chow, J.C., Doraiswamy, P. and Lowenthal, D.H., 2008. Source apportionment: findings from the US supersites program. Journal of the Air & Waste Management Association, 58(2), pp.265-288. https://doi.org/10.3155/1047-3289.58.2.265.

Wickham, H., 2014. Tidy Data. Journal of Statistical Software 59 (1), pp.1-23. https://doi.org/10.18637/jss.v059.i10.

Wickham, H., François, R., Henry, L. and Müller, K. 2021. dplyr: A Grammar of Data Manipulation, R package version 1.0.6, https://CRAN.R-project.org/package=dplyr.

Wood, S., 2017. Generalized Additive Models: An Introduction with R. 2nd edition. Chapman and Hall/CRC.

Wood, S., 2021. mgcv: Mixed GAM Computation Vehicle with Automatic Smoothness Estimation. R package version 1.8-38, URL https://CRAN.R-project.org/package=mgcv.

Zeileis, A., 2017. Modification suggested by the main author of strucchange to handle missing cases, https://stackoverflow.com/questions/43243548/strucchange-not-reporting-breakdates.

Zeileis, A., Kleiber, C. and Jackman, S., 2008. Regression Models for Count Data in R, Journal of Statistical Software, 27(8), pp. 1-25, https://doi.org/10.18637/jss.v027.i08.

Zeileis, A., Kleiber, C., Kraemer, W. and Hornik, K. 2003. Testing and Dating of Structural Changes in Practice, Computational Statistics & Data Analysis, 44, pp.109-123. https://doi.org/10.1016/S0167-9473(03)00030-6.

Zeileis, A., Leisch, F., Hornik, K. and Kleiber, C., 2002. strucchange: An R Package for Testing for Structural Change in Linear Regression Models, Journal of Statistical Software, 7(2), pp.1-38, http://www.jstatsoft.org/v07/i02/.

Supporting Information:

The results in this paper were obtained using R 4.1.3 and the AQEval 0.4.5 package. R itself and all packages used by AQEval and used in this article are available from the Comprehensive R Archive Network (CRAN) at https://CRAN.R-project.org/.

A. AQEval Installation Methods

From R, the most recent stable release versions of AQEval can be installed from CRAN using:

```
R> # Release version from CRAN
R> install.packages("AQEval")
```

The developers' versions of AQEval can be installed using:

```
R> # (if you do not have remotes package, install it from CRAN)
R> install.packages("remotes")

R> # developers version from Github
R> remotes::install_github("karlropkins/AQEval")
```

An AQEval package bundle (.zip or .tar.gz compressed file containing the package) can be installed using:

```
R> # developers version or other update as .tar.gz file (also using remotes)
R> remotes::install_local(file.choose()) # and select the bundle
```

B. Data Sources

Sources for datasets used in this article:

```
R> require(openair)
R> require(worldmet)
R> require(dplyr)

R> #1998-2005 data from London Marylebone Road AURN site
R> aq.my1 <- importAURN("my1", year=1998:2005) %>%
+         select(date, site, code, no2)

R> #1998-2005 data from Ealing Acton Town Hall KCL site
R> aq.ea2 <- importKCL("ea2", year=1998:2005, met=TRUE) %>%
+         select(date, site, code, no2, wd, ws)

R> #1998-2005 data from Kensington and Chelsea KCL site
R> aq.kc1 <- importKCL("kc1", year=1998:2005) %>%
+         select(date, site, code, no2)
```

C. AQEval Package Sources and Documentation

AQEval package versions:

- Stable release version on CRAN at https://CRAN.R-project.org/package=AQEval. And,
- Developers' version and code on GitHub at (https://github.com/karlropkins/AQEval).

Additional documentation:

- The project website is at https://karlropkins.github.io/AQEval/.

  Further documentation, guidance and news regarding package developments and updates can all be found here.