

This is a repository copy of *A trend analysis approach for air quality network data*.

White Rose Research Online URL for this paper:

<https://eprints.whiterose.ac.uk/id/eprint/145261/>

Version: Published Version

Article:

Lang, Polly E., Carslaw, David C. orcid.org/0000-0003-0991-950X and Moller, Sarah J. orcid.org/0000-0003-4923-9509 (2019) A trend analysis approach for air quality network data. *Atmospheric Environment: X*. 100030. ISSN: 2590-1621

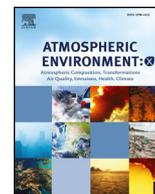
<https://doi.org/10.1016/j.aeaoa.2019.100030>

Reuse

This article is distributed under the terms of the Creative Commons Attribution-NonCommercial-NoDerivs (CC BY-NC-ND) licence. This licence only allows you to download this work and share it with others as long as you credit the authors, but you can't change the article in any way or use it commercially. More information and the full terms of the licence here: <https://creativecommons.org/licenses/>

Takedown

If you consider content in White Rose Research Online to be in breach of UK law, please notify us by emailing eprints@whiterose.ac.uk including the URL of the record and the reason for the withdrawal request.



A trend analysis approach for air quality network data

Polly E. Lang^{a,*}, David C. Carslaw^{a,c}, Sarah J. Moller^b

^a Wolfson Atmospheric Chemistry Laboratories, University of York, York, YO10 5DD, United Kingdom

^b National Centre for Atmospheric Science, Wolfson Atmospheric Chemistry Laboratories, University of York, York, YO10 5DD, United Kingdom

^c Ricardo Energy & Environment, Harwell, Oxfordshire, OX11 0QR, United Kingdom



ARTICLE INFO

Keywords:

Air quality
Monitoring network
Trend analysis
Bias
Rolling regression

ABSTRACT

Trend analysis of air pollutant concentrations becomes problematic when applied to data from air quality monitoring networks containing time series of differing lengths. The average trend from such data can be misleading due to biases in the monitoring network. For example, if new monitoring sites located in more polluted locations are added to a network, the introduction of these time series can leverage the trend upwards. A method for resolving this problem was developed, using rolling window regression to recursively calculate the change in pollutant concentration as a function of time, which can be used as a proxy for the true trend. The efficacy of the method was established by conducting simulations with known trends. The rolling change trend was shown to more accurately reflect the true trend than simply averaging the time series. Application of the technique to estimate trends in NO_x, NO₂ and NO₂/NO_x concentrations at London roadside monitoring sites over the period 2000–2017 revealed clear differences from the simple average. In particular, a significant monotonic downward trend in NO_x concentration was observed, in stark contrast to the average trend, which suggested little change in NO_x concentration had occurred over the same period. By accurately representing trends using time series of different lengths, this method has the benefit of being able to describe changes in air quality for locations and time periods with otherwise insufficient data.

1. Introduction

1.1. Background

Air quality monitoring networks are instrumental to the evaluation and management of air pollution by governments, policy makers and regulatory bodies. While other tools, such as emission inventories, are often used to track changes in emissions, the complex nature of atmospheric processes and local conditions mean that emissions data are not necessarily an accurate indicator of pollutant concentration or exposure. In contrast, ambient data from monitoring networks, subject to rigorous analysis, can reveal the pollutant concentrations, correlations and trends at measurement locations. Such information is invaluable for estimating the actual effects of social and infrastructure changes, and policy interventions on air quality.

Trend analysis is an important tool for examining the changes in pollutant concentration over time (Anttila and Tuovinen, 2010; Guerreiro et al., 2014), and can be used as evidence of the efficacy (or lack thereof) of policy interventions (Font and Fuller, 2016). In cases where the area under investigation contains a limited number of monitoring sites, a common approach is to compare the trends at each

individual monitoring site to yield an overall impression of changes in air quality in the area. For example, Mavroidis and Chaloulakou (2011) used this approach to estimate trends in particulate matter (PM) and ozone concentrations in Santiago, Chile 1989–1998 using data from four monitoring sites. The trends at each site were compared in order to establish a consensus, while differences between monitoring sites were rationalised using contextual information about each site. Some studies have attempted to replicate this approach with larger numbers of monitoring sites, such as the study by Masiol et al. (2017), which analysed the trends in concentration of a range of pollutants at 43 monitoring sites in the Veneto region of Italy. However, in the case of large regions or areas with an extensive monitoring network of sites available, this approach can be unwieldy and as such it may be beneficial to aggregate data from multiple monitoring sites to gain a representative view of the average air quality. Cluster analysis has been used to look at trends across a large number of sites allowing potential drivers for observed changes to be investigated and differences within and across regions to be explored (Malley et al., 2018).

Font and Fuller (2016) employed a different method to examine the trends in roadside increments of various pollutants between 2005–2009 and 2010–2014 by averaging data from 65 London monitoring sites.

* Corresponding author.

E-mail address: pl746@york.ac.uk (P.E. Lang).

Font and Fuller (2016) applied data capture filters and linear interpolation to ensure all individual time series from separate sites were of equal length. The problem with this approach is that data filtering excludes some information from the analysis. Fleming et al. (2018) in their analysis of ozone trends for the Total Ozone Assessment Report highlight that, particularly in developing countries, time series only span a few years and due to data capture requirements this limits the number of sites available for trend analysis. In this case the study is global and so there are still sufficient sites to provide the necessary data for robust trend analysis, but the distribution of the data across the globe is limited, with sparser sites in developing countries being more likely to be removed. For areas with sparser monitoring site coverage, or for trend analysis of long time periods, filtering the data may not be practicable, and therefore it may be necessary to average over all available monitoring sites to obtain a trend.

However, the trend in average concentration (average trend) over monitoring sites of differing duration is sensitive to biases in the monitoring network. Frequently, air pollution monitoring sites are moved to more polluted locations, closed in locations with low pollution levels, or new sites are opened in highly polluted locations that require more careful observation. The cumulative effect of site flux is often therefore that a monitoring network is increasingly biased towards monitoring sites with higher pollutant concentrations.

Duyzer et al. (2015) state that in their dual use for compliance monitoring and assessing population exposure, the choice of monitoring site location is made such as to provide data from the following: (i) the locations where the highest concentrations occur, and (ii) locations representative of the regional average. Typically, a distinction is made between roadside monitoring sites, which provide highly localised data from (i), and urban background monitoring sites, which are chosen to represent (ii). For this reason, movement of roadside monitoring sites to more polluted locations is not unexpected, but nonetheless has significant effects on the average trend. This issue was demonstrated in a 2014 report for the Department for Environment, Food & Rural Affairs (Defra, 2014). The long term trends in NO_2 and PM_{10} concentration were calculated using data from all monitoring sites in the AURN network, and compared to those derived using data from long term sites only. While the trends at urban background sites differed slightly, those from roadside sites displayed considerable differences, which were attributed to changes in monitoring site quantity and distribution over time.

In this paper, a range of techniques for identifying and mitigating the biasing effect of variation in time series length due to monitoring site flux on the average trend are developed. The efficacy and robustness of the methods are tested using simulated data. The methods are illustrated by a trend analysis of NO_x concentration, NO_2 concentration and NO_2/NO_x concentration ratio in London between 2000 and 2017 using data from roadside monitoring sites in the London air quality network. Finally, potential future applications of the new techniques are discussed.

London was chosen as the case study because of its unusual abundance of monitoring sites, including long term sites. However this situation is rare, giving rise to the need for methods that allow for the evaluation of the unbiased trend (i.e. the overall change in concentration across the network of monitoring sites) in the absence of long term monitoring sites.

2. Method

2.1. Identification of bias effects on the trend

Evidence of a bias in the monitoring network was sought by comparing trends averaged over (i) time series of differing lengths and (ii) time series of the same length. To this end, the trend in annual median concentration using data from (i) all monitoring sites and (ii) long term sites open for the entire duration of the period of study were compared.

In all cases, the average concentration was calculated using the median, as it is more robust to skewed data and the presence of extreme values.

This comparison is possible only if sufficient data is available from long term sites for the period of interest. In many cases, there may not be any reliable long term sites available as a basis of comparison. Additionally, any conclusions drawn from this comparison rely on the validity of the assumption that the trend from long term sites is representative of the true trend, and is not unduly affected by external influences.

In response to these limitations, a robust approach for observing and mitigating the effect of opening sites with high concentrations on the average trend was developed.

Rolling window regression (also known as rolling regression) is a technique most commonly used in time series analysis of financial data to examine variation in the output of a linear regression, such as the regression coefficient, over time (Wang and Zivot, 2006). The technique uses the same principle as a rolling average, except that a linear regression is applied to each time period (window) rather than an average. First, a rolling window width, n , is chosen. The data is partitioned into $N - n$ subsets, where N is the total number of observations in the time series. Each subset is rolled one observation ahead from the previous subset, resulting in a set of rolling windows of width n , each offset from the consecutive windows by one observation, and where the i th rolling window contains the observations $i, \dots, i + (n - 1)$. Linear regression is then applied to each rolling window.

A modification of traditional rolling regression was applied to the data, where each rolling window of width n contained data only from sites with measurements during every month within the period of the window (i.e. open and operational for all years within the window), ensuring that all time series within the window were of identical length.

Rolling trends in the concentration of the pollutant of interest for each window were plotted, resulting in a series of overlapping n year trends.

Comparison of the rolling trend and average trend using different values of n reveals a 'frame-by-frame' view of the potential bias. Each rolling trend overlaps with its neighbours for all years but one, and thus excludes data from monitoring sites opening in that year. In this way, by comparing trends in consecutive years, the effect of sites opening in that year can be visualised.

2.2. Extraction of the underlying trend

An optimal method to counter the influence of monitoring site flux on the average trend would aim to minimise the effect of the bias while retaining as much of the data as possible.

The simplest solution would be the exclusion of all sites not measuring constantly over the period of interest from the trend analysis via the application of a data capture filter. However, this approach would inevitably result in the sacrifice of a great deal of the available data, and in study areas with low numbers of sites could result in the conclusion that trend calculation was not possible. Furthermore, this method is predicated on the assumption that the long term sites are representative of the true trend in the location studied. Depending on the abundance (or lack thereof) of long-term sites, as well as other location-dependent external influences, this assumption may not be accurate.

An alternative method has been developed as an approach to this problem, with the advantage of retaining virtually all of the available data.

The method, which we shall refer to as the 'rolling change method', recursively calculates a concentration change, which approximates the trend in pollutant concentration. The concentration change in the first time point (e.g. the first year) is set as the median concentration over all monitoring sites in the first year. Next, the first moving window is defined as the period between time points $1, \dots, 1 + (n - 1)$. Data is drawn from the monitoring sites measuring throughout the duration of the window, and a linear regression is fit to the data, as described in

Section 2.1. The sum of the coefficient of the linear regression and the concentration change of the previous time point is assigned as the concentration change of the middle year of the moving window. The moving window is shifted down the time axis by one time unit (e.g. one year) and the process is repeated until the end of the time period of interest is reached.

For example, suppose the rolling change trend between 2000 and 2017 was calculated using a window width of three years. The starting point is the average of the annual average concentrations of all monitoring sites in 2000. The first moving window would select data from monitoring sites measuring constantly during 2000–2002, and fit a linear regression to the data. The sum of the regression coefficient and the concentration change in 2000 would be assigned as the concentration change for 2001. The moving window would then shift to 2001–2003 and repeat the process. The final moving window would use data from 2015 to 2017 to calculate the concentration change in 2016.

Similarly to the data filtering method described in Section 1.1, the rolling change method involves filtering monitoring sites by their data capture. However, unlike established methods, data filtering is applied over short windows of only 2–3 years, rather than the entire period of the trend analysis, therefore more data is retained during filtering.

Fig. 1 shows a schematic of the process, while Equations (1) and (2) describe the rolling regression and the recursive concentration calculation respectively. A more detailed algorithm can be found in Appendix A.

The terms in Equations (1) and (2) are defined as follows: Y_i is the variable of average concentration from all sites with sufficient data capture over the rolling window, i ; X_i is the variable of time points within the moving window, i ; x_i is the median year of X_i ; β_i is the coefficient of the rolling regression over the window, i ; ε_i is the irreducible error of the rolling regression, and Δy_i is the change in concentration assigned to the year x_i .

Equation (3) represents the rolling change trend itself. The trend is the concentration change (Δy_i) as a function of the median year of the rolling window (x_i).

$$Y_i = \beta_0 + \beta_i X_i + \varepsilon_i \quad (1)$$

$$\Delta y_i = \Delta y_{i-1} + \beta_i \quad (2)$$

$$\Delta y_i = f(x_i) + \varepsilon_i \quad (3)$$

The rolling change trend acts as a proxy for the trend in pollutant concentration, retaining information about the relative changes in concentration while discarding information regarding the relative magnitudes. The rolling change trend is constituted of rolling trends over n rolling windows, each fit to a set of time series of identical length. In this way, the leveraging effect induced by the inclusion of high magnitude time series does not affect the trend, so data from all sites with a duration of at least n years can be included in the analysis. The choice of n dictates the criteria for inclusion of monitoring sites into the analysis. Larger values of n impose more stringent requirements for site duration, and thus exclude more monitoring sites.

The functions used for the trend analysis in the paper, including the calculation of rolling trends and rolling change trends, are available in the aqtrends R package (Lang, 2018).

2.3. Description of data

The data used in the London case study were sourced from the Automatic Urban and Rural Network (AURN) maintained by Defra, the London Air Quality Network (LAQN) run by King's College London, and the Air Quality England database collected by Ricardo Energy & Environment.

Each of these networks contains a number of monitoring sites, which record hourly observations of air pollutant concentrations. The concentrations of NO_x and NO_2 were measured using the European

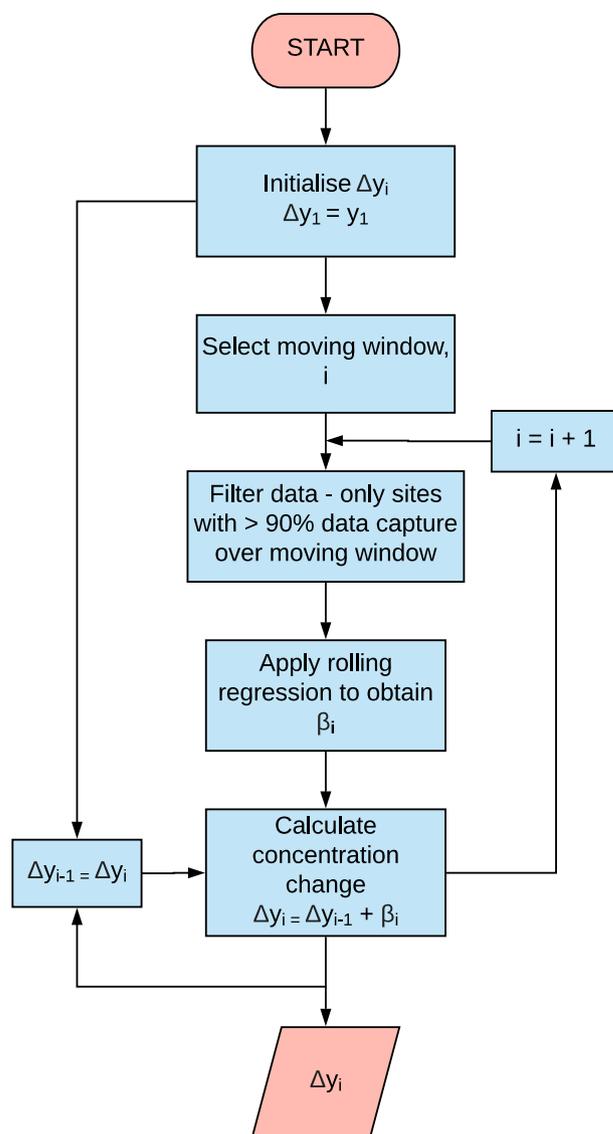


Fig. 1. Schematic of the rolling change method. The output for the process as a whole (the concentration change for the rolling window), i , is highlighted in red. (For interpretation of the references to colour in this figure legend, the reader is referred to the Web version of this article.)

Commission reference method of chemiluminescence with molybdenum converter.

For each monitoring site, data more than 10 times the interquartile range from the upper quartile was considered to be an outlier and removed from the data set. Monitoring sites with less than 75% data capture over the period during which they were measuring data were not included in the trend analysis. The mean and the standard deviation of the hourly NO_x and NO_2 concentrations measured at each monitoring site is given in Supplementary Material 1.

The hourly data was used to calculate annual average concentrations using three different methods of trend analysis. For the average trend (all sites), all available data from all monitoring sites measuring during the period of analysis was included in the average (median). The average trend (using data from long term sites only) was calculated using data only from sites measuring throughout the duration of the period of analysis. This was defined as recording measurements during every month within the period of analysis. Additionally, a data capture criterion was applied to ensure that all long term sites had at least 90% data capture over the period of analysis. Finally, for the rolling change method, within each moving window, only data from sites with

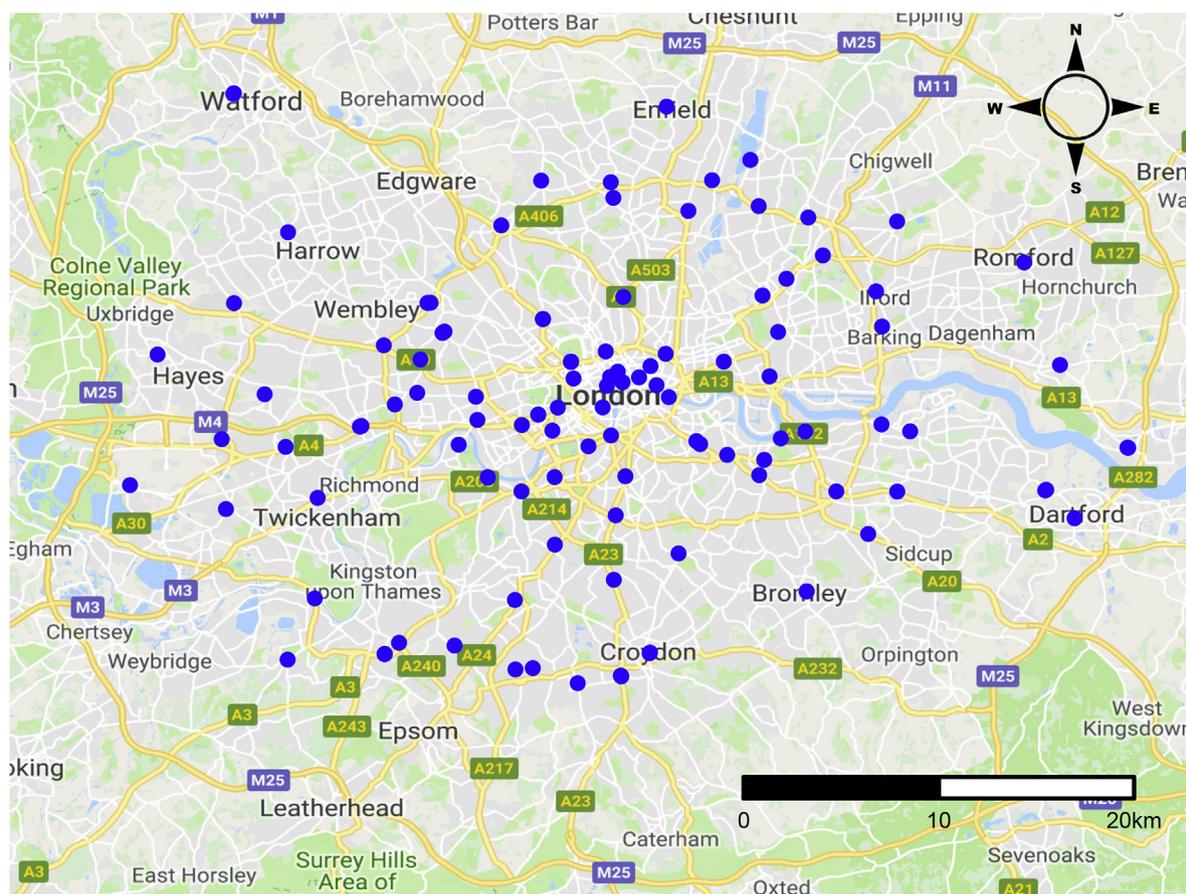


Fig. 2. Map showing the locations of the London roadside monitoring sites measuring NO_x and NO_2 used in the analysis. More information on individual sites can be found in Supplementary Material 1.

measurements during every month within the period of the window was included in the calculation for that window.

London monitoring sites were selected as all sites within a bounding box of coordinates 51.25°N , 51.71°N , -0.54°E , 0.28°E . This box was roughly equivalent to the boundary of the M25 orbital motorway. The data included 121 roadside sites and 99 urban background sites measuring over the period 2000–2017. The roadside monitoring sites are shown in Fig. 2. Of the 121 roadside sites, 102 sites measuring NO_x and 105 sites measuring NO_2 met the data capture requirements for the trend analysis. Of these, 9 sites measuring NO_x and 10 measuring NO_2 were open for the entire duration of the period of trend analysis (long term sites). More information about individual sites is given in Supplementary Material 1.

All data importing, cleaning, transformation and analysis was carried out in R.

3. Results and discussion

3.1. Testing the rolling change method through simulations

Simulations were carried out to compare the effectiveness of the average trend and the rolling change trend to display the true change in pollutant concentration over time. Data were simulated to mimic the properties of the real monitoring data, but with the true trend known. To reflect the various possible properties of a monitoring network, data were generated from four scenarios:

(a) Long term monitoring sites. All of the time series had the same true trend (with noise added) and the same length (equal to the length of the entire time period (2000–2017)). Variation in the

concentrations of different time series was simulated by sampling the concentration in the first year of the time series from a normal distribution with a mean equal to the concentration of the true trend in that year and a standard deviation of 10 ($X \sim \mathcal{N}(\text{true trend concentration}, 10^2)$).

- (b) Short term monitoring sites without a time-dependent bias in concentration. All of the time series had the same underlying trend, but different lengths. The starting year of each time series was randomly sampled from the standard uniform distribution, constrained between 2000 and 2015. The time series length was also randomly sampled from the standard uniform distribution, $U(0,1)$. Variation in the concentration of different time series was simulated using the same method as described in (a) above.
- (c) Short term monitoring sites with a time-dependent bias. Each time series had the same underlying trend, but different lengths. Additionally, in order to simulate the effect of increasing bias towards more polluted locations over time, the simulated concentration in the first year of the time series was randomly drawn from the standard uniform distribution, and multiplied by a bias factor proportional to the starting year of the time series. The result was that the concentration in latter years was more likely to be higher than in former years. The bias factor took the form $y_i = 1 + 0.08x_i + \varepsilon_i$ where y was the value of the bias factor, x was the index of the starting year of the time series (between 1 and 18), and ε was the random error. The error for each value of the bias factor, ε_i , was randomly sampled from the normal distribution $N(0,0.5)$.
- (d) A combination of time series generated according to the ‘long term’ scenario and the ‘short term with bias’ scenario. The method of generating each time series was determined by random selection,

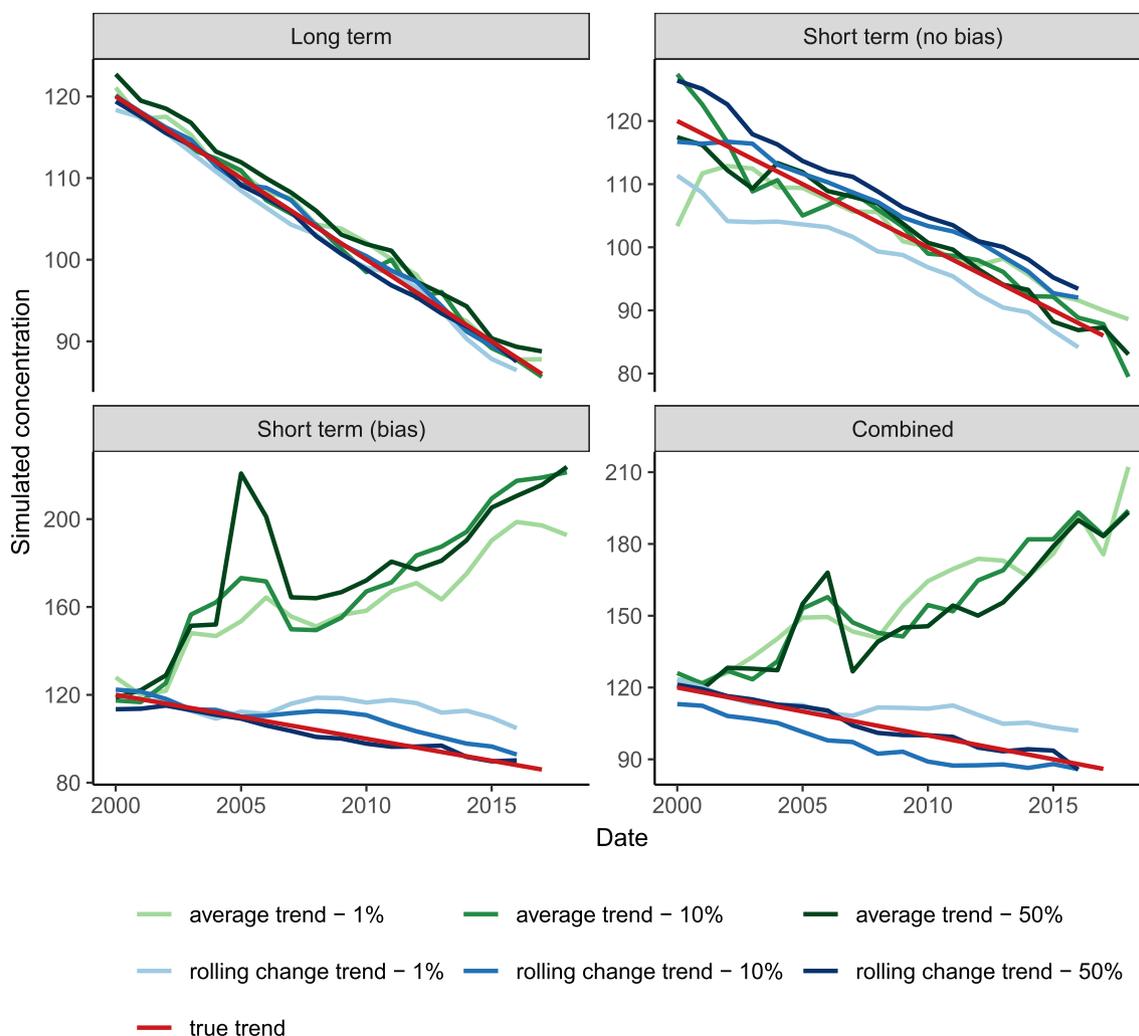


Fig. 3. Comparison of the average trend and rolling change trend ($n = 3$) with the true trend of simulated data for four different scenarios. In each case, the trends are derived from 100 random samples, each of 100 simulated time series. The lines correspond to the trends with NCC equal to the 50th, 10th and 1st percentile of the NCC distribution over all 100 sampled trends — in other words, the median trend, the 10th worst trend and the worst trend, with respect to the similarity to the true trend.

where the probability of generating a short term site was ten times as likely as that of generating a long term site, in line with the observed proportions of long term and short term sites in the London roadside monitoring network.

For each scenario, 100 sets of simulated data, each consisting of 100 simulated time series, were randomly sampled. The rolling change trend and the average trend were calculated for each sample of simulated data, and their similarity to the ‘true trend’ (the function used to create the simulated data) was evaluated using normalised cross-correlation (NCC). The normalised cross-correlation of two time series is a value between 1 and -1, where 1 means the two time series are perfectly correlated, while -1 corresponds to perfect anti-correlation. The results are shown in Fig. 3.

The average trend in Scenarios (c) and (d) was considerably biased relative to the known true trend, as was observed in the real data, but in each case the rolling change trend provided a more accurate representation of the true trend.

Furthermore, the slope of the rolling change trend was shown to be more accurate than that of the average trend. The slopes of each sampled rolling change trend and average trend were calculated using the Theil-Sen estimator, and compared to the slope of the true trend from which the data were simulated to derive the percentage error. For the

combined scenario, the median error of the rolling change trend was 15%, while for the average trend the error was 293%.

The suitability of the technique for situations with limited data available was also evaluated by applying the trend analyses to 100 samples of 4 time series simulated using the ‘combined’ scenario, as shown in Fig. 4. As before, the rolling change trend represented the true trend with greater accuracy than the average trend, indicating that the method extends well to situations with a very limited number of monitoring sites.

Simulated data was also used to demonstrate that the rolling change method is robust to the use of different values of the moving window width, n , as shown in Supplementary Material 2. The accuracy of the rolling change method increases slightly as the window width increases, however the amount of data filtered out also increases. To achieve a reasonable balance between maximising the accuracy of the rolling change trend, while maximising the amount of data retained in the analysis, a window width of $n = 3$ was used in the following applications of the method.

3.2. Long term trends in London ambient air quality

3.2.1. Identification of the bias effect on the trend

Comparison of the average trend over all London roadside sites

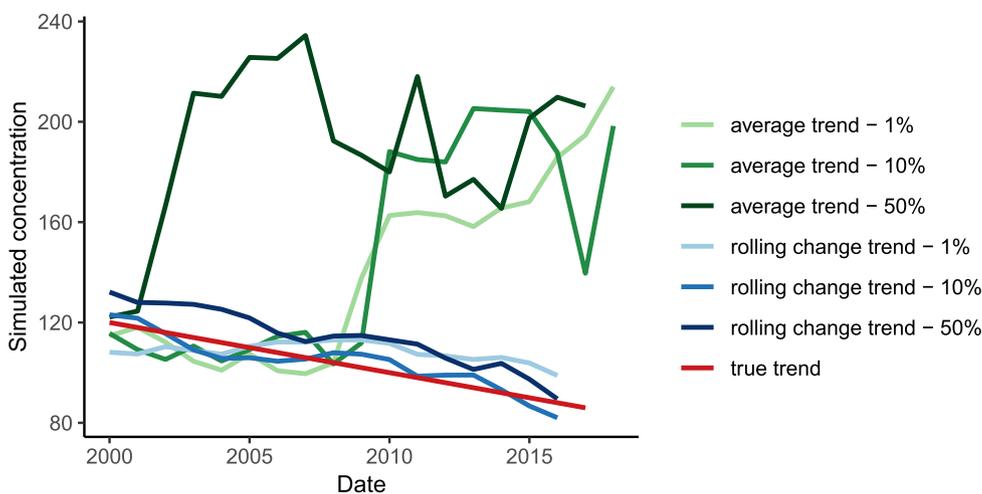


Fig. 4. Comparison of the average trend and rolling change trend ($n = 3$) with the true trend of data from 5 time series simulated using the ‘combined’ scenario. The trends are derived from 100 random samples of simulated data. The lines correspond to the trends with NCC equal to the 50th, 10th and 1st percentile of the NCC distribution over all 100 sampled trends — in other words, the median trend, the 10th worst trend and the worst trend, with respect to the similarity to the true trend.

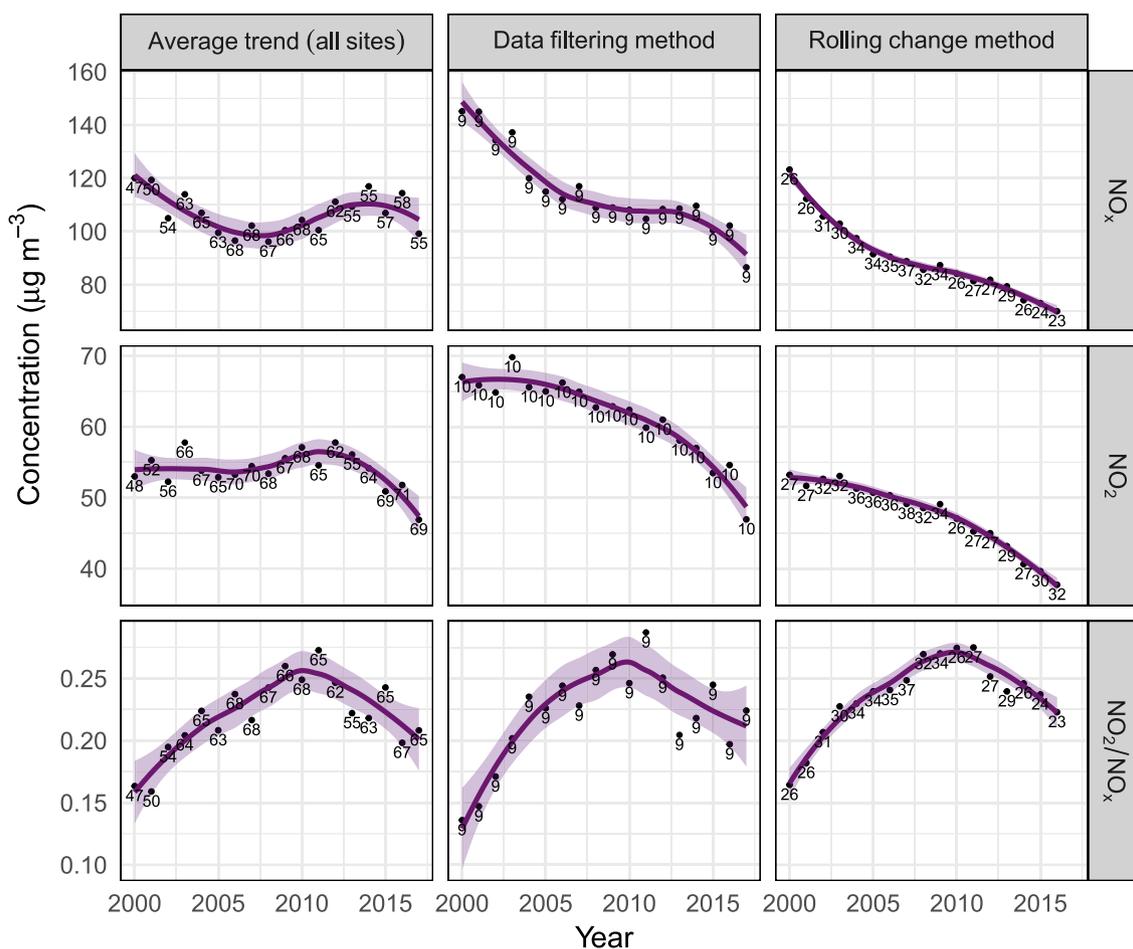


Fig. 5. Comparison of the rolling change trends in NO_x concentration, NO_2 concentration, and NO_2/NO_x ratio at London roadside sites 2000–2017, using $n = 3$ (‘Rolling change method’) with the trend in the average concentration using data from (i) all available monitoring sites (‘Average trend (all sites)’) and (ii) long term sites only (‘Data filtering method’). The lines represent a loess smooth fit to the data, and the shaded bands represent the 95% confidence interval around the smooth. The numbers at each data point correspond to the number of monitoring sites contributing to the data point.

during the period 2000–2017 with the average trend over long term sites (those measuring constantly over the same time period) reveals a dramatic difference in trend, as shown in the two left-hand plots in Fig. 5. The trend of the long term sites is constituted of data from between nine and eleven monitoring sites. Therefore the disparity is unlikely to be the result of lack of representativeness due to local influences. A more likely explanation is a bias towards opening new monitoring sites in increasingly polluted locations, resulting in the

sudden introduction of high concentration time series causing abrupt increases in the average concentration despite no commensurate increase in the trends at individual sites.

The increase in bias in site location towards more polluted sites over time was affirmed by comparing the median annual ambient concentrations at roadside monitoring sites opening and sites closing in a given year across the period studied (see Supplementary Material 3). The difference between the average concentration at sites that are

opening and those that are closing is positive (i.e. concentrations are higher at sites that are opening) over almost all years for NO_x and NO₂.

The effect of the bias in site location on the trend in average roadside NO_x and NO₂ concentrations can be observed through a comparison of the rolling trends over rolling windows of different widths (*n*), as shown in Supplementary Material 4.

When the same trend analysis was applied to monitoring data from London urban background sites, however, no bias in the average trend was observed (see Supplementary Material 5), in corroboration of the findings of the Defra report discussed in Section 1.1 (Defra, 2014). This is, in part, because the bias towards opening sites in more polluted locations is far less pronounced for urban background sites, which also move less frequently than do roadside sites. Moreover, any bias in site location is likely to have a smaller effect on the average trend at urban background sites, because the NO_x and NO₂ concentrations are dominated by non-local background sources rather than local traffic sources, which constitute the major source at roadside sites.

3.2.2. Extraction of the underlying trend

Having established the existence of a bias effect on the average trend by the short term sites, the next step is to mitigate this bias effect in order to reveal the true underlying trend. The rolling change method described in Section 2.2 was applied to the London roadside monitoring data.

The rolling change trends in NO_x concentration, NO₂ concentration and NO₂/NO_x ratio are shown in Fig. 5 (right). In all cases, these derived trends bear a far closer similarity with the trend for the long-term site (Fig. 5 (middle)) than with the biased average trend (Fig. 5 (left)), offering further evidence in support of the technique's efficacy.

The rolling change technique reveals a more optimistic trend from 2000 to 2017 in NO_x concentration at London roadside sites than that implied by the average trend. Table 1 shows the Theil-Sen slopes of the trends derived using the three different methods (the trend in average concentration using data from (i) all sites and (ii) long term sites only, and the trend derived using the rolling change method).

Application of the Theil-Sen estimator to the NO_x concentration trends in Fig. 5 yielded a slope of -2.52 [-3.32 , -1.96] g m⁻³ year⁻¹ for the rolling change trend. In contrast, the gradient of the average trend was -0.22 [-1.23 , 1.00] g m⁻³ year⁻¹. The rolling change trend is a highly monotonic, almost linear decrease, while the average trend indicates a fluctuation with initial decrease to 2007, followed by a period of increase to 2013–14, with little overall change in NO_x concentration.

The differences between the average and rolling change trends in NO₂ concentration were less extreme, but nonetheless notable. Theil-Sen slope of the rolling change trend was -0.90 [-1.13 , -0.69] g m⁻³ year⁻¹ in comparison to -0.08 [-0.46 , 0.24] g m⁻³ year⁻¹ for the average trend. The rolling change trend revealed a monotonic downwards trend since 2003–4, with an increasingly steep gradient in later years, while the average trend does not show any downward inclination until 2012–13, and even shows a slight increase between 2008 and

2012.

The effectiveness with which the rolling change trend represents the 'true trend' was evaluated by comparison with trends in NO_x and NO₂ from emissions data, satellite data and previous studies of London air quality.

The rolling change trend incorporates information from more monitoring sites than would be possible using only long term sites or individual sites. As such, it is more likely to be reflective of overall trends in traffic emissions across London, and therefore more comparable with trends estimated by emissions inventories. The UK trend in NO_x emissions from urban driving sources (NAEI, 2018) is shown in Fig. 6. The emissions data shows a monotonic, almost linear downward trend between 2000 and 2016, similar to the rolling change trend in NO_x concentration from the London data (see Fig. 5). The emissions trend shows a -56% change from 2000 to 2016, which is not dissimilar to the -43% change in the rolling change trend in NO_x concentration over the same period. A smaller slope is expected for the ambient concentration trend than the emissions trend because concentrations at roadside are dominated by traffic sources but other sources also contribute. One such source is natural gas combustion for domestic heating, from which NO_x emissions have decreased less between 2000 and 2017 than emissions from road transport sources, effectively depressing the slope of the trend in ambient NO_x concentration relative to the trend in NO_x emissions from transport sources (NAEI, 2018; Wakeling et al., 2018).

A recent study of London air quality using satellite data estimated a trend in NO₂ concentration of -0.23×10^5 molecules cm⁻² year⁻¹ between 2005 and 2015, which is approximately -1.76 g m⁻³ year⁻¹, assuming a column height of 10 km (Pope et al., 2018). The slope of the rolling change trend (with 95% confidence intervals given in brackets) in NO₂ concentration over the same period from roadside monitoring sites was -1.03 [-1.48 , -0.74] g m⁻³ year⁻¹, compared to the average trend slope of -0.02 [-0.41 , 0.31] g m⁻³ year⁻¹. While neither trend indicates as large a downward trend as that from the satellite data, the rolling change trend provides concordant evidence of a negative trend in NO₂ over this period. Some disparity between the satellite data and monitoring data is expected, because the satellite measurements integrate concentrations across the entirety of London, while the ambient concentration data were measured exclusively at roadside monitoring sites. As a result, the long term trends in the satellite data will be driven by multiple sources, including domestic activity and power station emissions, in contrast to the trends in ambient concentration which are heavily dominated by traffic sources.

The rolling change trends also corroborate the findings of Grange et al. (2017) that the NO₂/NO_x vehicle emission ratio across Europe follows a pattern of increase from 1995 to 2008 then decrease between 2009 and 2015. This pattern is replicated in the NO₂/NO_x rolling change trend shown in Fig. 5 and reflects changes to the direct emission of NO₂ from diesel vehicles.

A comparison of the results of the study by Font and Fuller (2016) examining trends in roadside increments of NO_x and NO₂ concentration

Table 1

Theil-Sen slope and 95% confidence intervals of the trend in average concentration (all sites), the trend in average concentration (long term sites) and the rolling change trend in NO_x, NO₂ and NO₂/NO_x concentration at roadside in London 2000–2017.

Pollutant	Method	Theil-Sen slope (g m ⁻³ year ⁻¹)	95% confidence interval
NO _x	Average trend (all sites)	-0.22	[-1.23, 1.00]
NO _x	Average trend (longterm sites)	-2.59	[-3.33, -1.37]
NO _x	Rolling change method	-2.52	[-3.32, -1.96]
NO ₂	Average trend (all sites)	-0.08	[-0.46, 0.24]
NO ₂	Average trend (longterm sites)	-0.95	[-1.19, -0.62]
NO ₂	Rolling change method	-0.90	[-1.13, -0.69]
NO ₂ /NO _x	Average trend (all sites)	0.00	[-0.00, 0.01]
NO ₂ /NO _x	Average trend (longterm sites)	0.00	[-0.00, 0.01]
NO ₂ /NO _x	Rolling change method	0.00	[-0.00, 0.01]

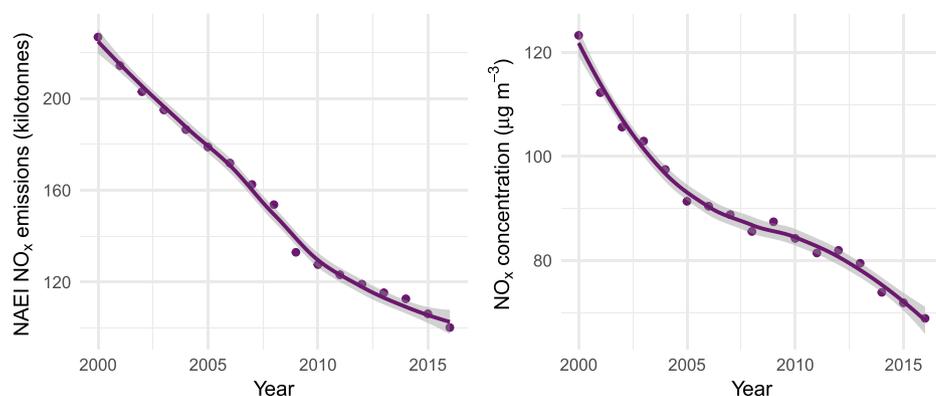


Fig. 6. Trend in UK NO_x emissions from road transport (urban driving) sectors between 2000 and 2016 (left) compared to the rolling change trend in NO_x concentration over the same period (right). The lines represent a loess smooth fit to the data, and the shaded bands represent the 95% confidence interval around the smooth.

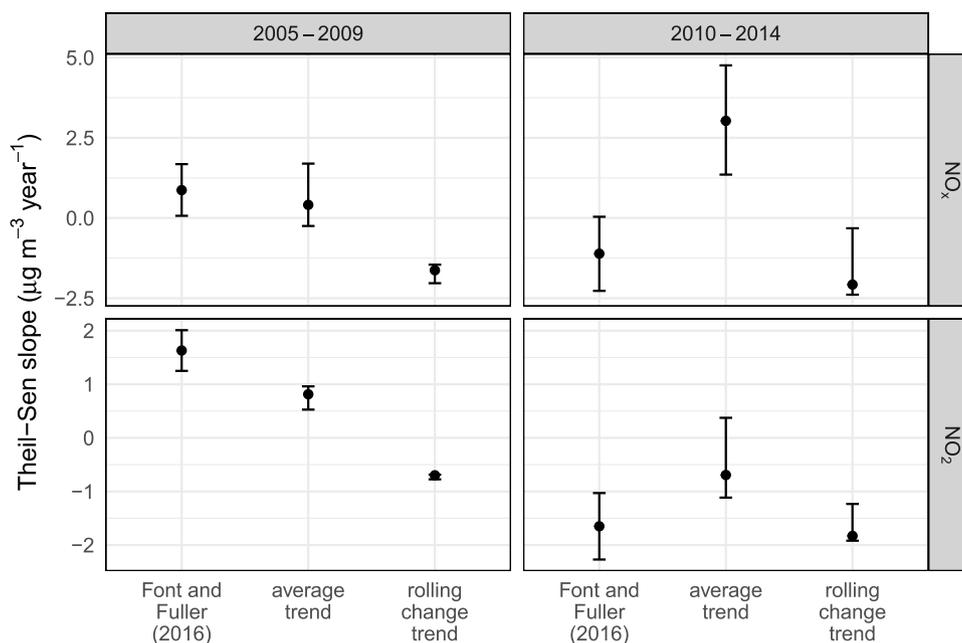


Fig. 7. Comparison of the Theil-Sen slope calculated by Font and Fuller (2016) with the rolling change trend and the average trend in NO_x and NO₂ roadside increments at London roadside monitoring sites between 2005–2009 and 2010–2014. The error bars represent 95% confidence intervals.

in London between 2005–2009 and 2010–2014 with those obtained from the rolling change trend and the average trend are shown in Fig. 7. As mentioned in Section 1.1, Font and Fuller (2016) applied data capture filters and linear interpolation to include only time series of similar length in the analysis. As a result, some data were excluded, leaving data from 47 monitoring sites from which to derive trends. In contrast, the use of the rolling change technique allowed for inclusion of data from all available monitoring sites, which for 2005–2009 was 91 and 93 sites for NO_x and NO₂ respectively, and for 2010–2014, 85 and 86 sites respectively.

As can be seen in Fig. 7, for the period 2010–2014, the slope of the rolling change trend was more similar to the trend calculated by Font and Fuller (2016) than that of the average trend, although for the period 2005–2009, the rolling change trend differed considerably from that calculated by Font and Fuller (2016). Positive trends were observed for both NO_x and NO₂ between 2005 and 2009 by Font and Fuller (2016), while negative trends were observed using the rolling change method. However, negative slopes were observed for both NO_x and NO₂ concentrations between 2010 and 2014, in corroboration of the findings of Font and Fuller (2016).

Font and Fuller (2016) took advantage of the unusual abundance of monitoring sites in London to implement a filtering method while

retaining enough data to robustly represent the overall trend in concentration. However, the applicability of this approach is limited to situations with a similar abundance of monitoring sites available, excluding most urban areas. In these cases, the rolling trend method may be the only robust method of calculating an overall long term trend in ambient concentration.

Additionally, the data filtering method implemented by Font and Fuller (2016) limits the time period over which the long term trend can be analysed to periods over which a sufficient number of monitoring sites are measuring constantly. For example, in an eighteen year trend analysis of NO_x or NO₂ concentrations, such as the one demonstrated in Section 3, the application of the data filtering method would constrain the analysis to data from only nine or ten monitoring sites. In other locations, it is unlikely that any monitoring sites have been measuring constantly for eighteen years, and such a long term analysis would be impossible.

Finally, as alluded to previously, data filtering methods are wasteful. By excluding monitoring sites which are not measuring constantly over the period of interest, a great deal of potentially important data is not considered. The rolling change method's advantage over traditional techniques is that it does not automatically exclude data from short term monitoring sites, and so retains far more of the data in the

analysis.

3.3. Potential applications

The rolling change method offers the following advantages over traditional methods of trend analysis:

- Robust long term trend analysis across monitoring networks which may be subject to time-dependent biases
- Enables long term trend analysis to be undertaken for areas with few/no long term monitoring sites

A lack of long term roadside monitoring sites is a major barrier to the analysis of long term trends in roadside pollutant concentrations. As previously mentioned in Section 1.1, roadside monitoring sites are frequently re-located to locations deemed more critical for compliance monitoring, resulting in short time series. To illustrate this difficulty, suppose the trend analysis of roadside NO_x and NO₂ concentrations between 2000 and 2017 was carried out for other UK cities (excluding London). In the UK, there are 4 functional urban areas (FUA) and 4 towns (excluding London) with long term roadside monitoring sites measuring NO_x and NO₂ concentration over the period 2000–2017, none of which has more than 1 monitoring sites. The scarcity of long term roadside monitoring sites poses a serious problem for comprehensive long term trend analysis. However, use of the rolling change method allows the relaxation of the constraint limiting the useable data to that from long term monitoring sites. As a consequence, the range of locations in which long term trend analysis is possible can be expanded to areas which would be inaccessible using the established methods, such as data filtering.

Moreover, the technique is broadly applicable to any situation requiring the aggregation of multiple, concurrent time series of differing lengths into a single, overall trend. For example, such a situation may arise in other environmental sciences where continuous monitoring is carried out over a network of sites, such as water quality monitoring, soil monitoring or oceanography.

Even outside the environmental sciences, trend analysis of multiple time series is routinely carried out in finance, quality control and the social sciences. In these fields, as in environmental monitoring, it is more usual for analysis to be limited to time series of the same length. However, with the rapid growth of sensor technologies and the commensurate increase in the automatic collection of time series data, the ability to analyse variable length time series could be advantageous.

4. Conclusions

Air quality monitoring networks offer the potential to visualise and quantify long-term trends over large regions through aggregation of data from multiple monitoring sites. However, analysis of roadside monitoring site data from the London network suggests caution is required when averaging data from a monitoring network containing time series of variable duration. Movement, opening and closing of monitoring sites introduce biases into the average trend, resulting in a misleading view of the changes in air quality.

Techniques were developed with the aim of identifying and mitigating these influences to robustly represent the true long term trend. In particular, a method involving the calculation of a change in concentration using rolling window regression was developed as an effective alternative to simple averaging. This technique was demonstrated to estimate the true trend in pollutant concentration with far greater accuracy than the simple average trend when applied to a set of time series of disparate lengths.

The ability to use multiple time series of differing lengths in trend analysis offers potential advantages for air quality and environmental monitoring applications, as well as time series analysis in other fields. An important advantage of the technique is that it maximises the use of

the information available and is suited to situations where a large number of monitoring sites may not be available but where an aggregate view of overall changes in concentrations is still valuable.

Acknowledgements

Polly E. Lang was funded by the Department of Chemistry at the University of York. The authors would like to acknowledge Defra, King's College London and Ricardo Energy & Environment for providing the data used in this study.

Appendix A. Supplementary data

Supplementary data to this article can be found online at <https://doi.org/10.1016/j.aeoa.2019.100030>.

Appendix A

The algorithm for the rolling change method described in Section 2.2 is as follows:

1. Choose the time range over which to calculate the trend, and the value of the rolling window width, n .
2. Initialise Δy_i as the average of the annual average concentrations of all monitoring sites in the first year, y_i .
3. Identify the moving window, i , as the time period $x_i, \dots, x_{i+(n-1)}$.
4. Select the vector of dates encapsulated by the moving window, X_i .
5. Filter the concentration data to include only data from sites with $\geq 90\%$ data capture over the moving window. The result will be a vector of concentration values of length n , Y_i .
6. Fit a linear regression model to the filtered concentration data, (X_i, Y_i) , as in Equation (1).
7. Calculate the concentration change over the moving window using the regression coefficient, β_i , and the concentration change of the previous window, Δy_{i-1} using Equation (2).
8. Assign the concentration change, Δy_i , to the median date of the rolling window, x_i .
9. Slide the moving window by one time point towards the end of the time range.
10. Repeat Steps 4–9 until the moving window reaches the end of the time range.
11. The rolling change trend is Δy_i as a function of x_i over all i (i.e. the entire time range), as shown in Equation (3).

References

- Anttila, P., Tuovinen, J.-P., Jan 2010. Trends of primary and secondary pollutant concentrations in Finland in 1994/2007. *Atmos. Environ.* 44 (1), 30–41.
- Defra, 2014. Air Pollution in the UK 2013. (Tech. rep).
- Duyzer, J., van den Hout, D., Zandveld, P., van Ratingen, S., Mar 2015. Representativeness of air quality monitoring networks. *Atmos. Environ.* 104, 88–101.
- Fleming, Z.L., Doherty, R.M., Von Schneidmesser, E., Malley, C.S., Cooper, O.R., Pinto, J.P., Colette, A., Xu, X., Simpson, D., Schultz, M.G., Lefohn, A.S., Hamad, S., Moolla, R., Solberg, S., Feng, Z., Feb 2018. Tropospheric Ozone Assessment Report: present-day ozone distribution and trends relevant to human health. *Elem Sci Anth* 6 (1), 12.
- Font, A., Fuller, G.W., Nov 2016. Did policies to abate atmospheric emissions from traffic have a positive effect in London? *Environ. Pollut.* 218, 463–474.
- Grange, S.K., Lewis, A.C., Moller, S.J., Carslaw, D.C., Dec 2017. Lower vehicular primary emissions of NO₂ in Europe than assumed in policy projections. *Nat. Geosci.* 10 (12), 914–918.
- Guerreiro, C.B., Foltescu, V., de Leeuw, F., Dec 2014. Air quality status and trends in Europe. *Atmos. Environ.* 98, 376–384.
- Lang, P., 2018. Aqtrends. <https://github.com/pollylang/aqtrends>.
- Malley, C.S., Von Schneidmesser, E., Moller, S., Braban, C.F., Hicks, W.K., Heal, M.R., 2018. Analysis of the distributions of hourly NO₂ concentrations contributing to annual average NO₂ concentrations across the European monitoring network between 2000 and 2014. *Atmos. Chem. Phys.* 18, 3563–3587.
- Masiol, M., Squizzato, S., Formenton, G., Harrison, R.M., Agostinelli, C., Jan 2017. Air quality across a European hotspot: spatial gradients, seasonality, diurnal cycles and trends in the Veneto region, NE Italy. *Sci. Total Environ.* 576, 210–224.
- Mavroidis, I., Chaloulakou, A., Dec 2011. Long-term trends of primary and secondary NO₂ production in the Athens area. Variation of the NO₂/NO_x ratio. *Atmos. Environ.* 45 (38), 6872–6879.

- NAEI, 2018. UK National Atmospheric Emissions Inventory. <http://naei.beis.gov.uk/index>.
- Pope, R.J., Arnold, S.R., Chipperfield, M.P., Latter, B.G., Siddans, R., Kerridge, B.J., Jun 2018. Widespread changes in UK air quality observed from space. *Atmos. Sci. Lett.* 19 (5), 1–8.
- Wakeling, D., Passant, N., Murrells, T., Misra, A., Pang, Y., Thistlethwaite, G., Walker, C., Brown, P., Del Vento, S., Hunter, R., Wiltshire, J., Broomfield, M., Watterson, J., Pearson, B., Rushton, K., 2018. UK Informative Inventory Report (1990 to 2016) (Tech. rep).
- Wang, J., Zivot, E., 2006. Rolling analysis of time series. In: *Modeling Financial Time Series with S-PLUS*. Springer New York, New York, NY, pp. 313–360 Ch. 9.