Article:

Grange, Stuart K. orcid.org/0000-0003-4093-3596 and Carslaw, David C. orcid.org/0000-0003-0991-950X (2018) Using meteorological normalisation to detect interventions in air quality time series. Science of The Total Environment, 653. pp. 578-588.

Using meteorological normalisation to detect interventions in air quality time series

Stuart K. Grange[a,*], David C. Carslaw[a,b]

[a] *Wolfson Atmospheric Chemistry Laboratories, University of York, York, YO10 5DD, United Kingdom*
[b] *Ricardo Energy & Environment, Harwell, Oxfordshire, OX11 0QR, United Kingdom*

Abstract

Interventions used to improve air quality are often difficult to detect in air quality time series due to the complex nature of the atmosphere. Meteorological normalisation is a technique which controls for meteorology/weather over time in an air quality time series so intervention exploration (and trend analysis) can be assessed in a robust way. A meteorological normalisation technique, based on the random forest machine learning algorithm was applied to routinely collected observations from two locations where known interventions were imposed on transportation activities which were expected to change ambient pollutant concentrations. The application of progressively stringent limits on the content of sulfur in marine fuels was very clearly represented in ambient sulfur dioxide (SO_2) monitoring data in Dover, a port city in the South East of England. When the technique was applied to the oxides of nitrogen (NO_x and NO_2) time series at London Marylebone Road (a Central London monitoring site located in a complex urban environment), the normalised time series highlighted clear changes in NO_2 and NO_x which were linked to changes in primary (directly emitted) NO_2 emissions at the location. The clear features in the time series were illuminated by the meteorological normalisation procedure and were not observable in the raw concentration data alone. The lack of a need for specialised inputs, and the efficient handling of collinearity and interaction effects makes the technique flexible and suitable for a range of potential applications for air quality intervention exploration.

Keywords:
Air pollution, Data analysis, Management, Machine learning, Random forest

1. Introduction

Across all spatial and temporal scales, weather influences concentrations of atmospheric pollutants and in turn ambient air quality (Stull, 1988; Monks et al., 2009). The effects of weather (or meteorology) on air quality are often much greater than intervention or management efforts to control air pollution and therefore intervention events can be very difficult to detect and quantify within an observational record (Anh et al., 1997). Similarly, when considering trends in ambient air pollution, it can be difficult to know whether a change in concentration is due to meteorology or a change in emission source strength. Meteorological variation can therefore frustrate the analysis of trends in different pollutant species. If meteorology is not controlled or accounted for, the changes in pollutant concentrations observed may be contaminated with meteorological variation rather than emission or chemically induced perturbations which can lead to erroneous conclusions concerning the efficacy of air quality management strategies (Libiseller et al., 2005; Wise and Comrie, 2005). This issue is often acknowledged, but infrequently addressed.

Meteorological normalisation is one technique which can be used to control for meteorology over time in air quality time series. The central philosophy of meteorological normalisation is to reduce variability in an air quality time series with statistical modelling. The reduction of variability is achieved by training a model which can explain some of the variation of pollutant concentrations through a number of independent variables. The independent variables used are typically surface-based meteorological observations and time variables which act as proxies for regular emission patterns such as hour of day and season (Derwent et al., 1995). However, in practice, any independent variable which could explain variations in pollutant concentrations could be used. Once the model has been trained and it is found that it can explain an adequate amount of the dependent variable's variation, the model can be used to remove the influence the independent variables have on the dependent variable by sampling and predicting. The time series which results can then be exposed to further exploratory data analysis (EDA) techniques such as formal trend analysis and/or intervention

*Corresponding author
Email address: stuart.grange@york.ac.uk (Stuart K. Grange)

exploration (Grange et al., 2018). The normalised time series is in the pollutant's original units and can be thought of as concentrations in "average" or invariant weather conditions.

There has been some air quality research conducted which uses the idea of change-point analysis to investigate changes in atmospheric pollutant concentrations (for example Carslaw et al., 2006; Carslaw and Carslaw, 2007). Methods such as these rely on regime changes where a time series abruptly shifts from one regime to another (Lyubchich et al., 2013). In the air quality domain, this rarely happens, since changes are usually nuanced and occur progressively with much variability which makes the generality of this approach for investigating intervention efforts poor. Meteorological normalisation is potentially a more general approach which enables its use in a greater range of applications.

Atmospheric processes are complex, non-linear, and observations commonly record collinearity with other observations. These attributes make the process of statistical modelling very challenging, especially so with parametric methods (Barmpadimos et al., 2011). With the rise of machine learning algorithms, these attributes can be much more easily accommodated due to the non-parametric and robust nature of these techniques (Friedman et al., 2001). The meteorological normalisation technique used here uses random forest, an ensemble decision tree machine learning method as the modelling algorithm.

Random forest has been described very well and in depth elsewhere (see Breiman, 2001; Friedman et al., 2001; Tong et al., 2003; Ziegler and König, 2013; Jones and Linder, 2015; Grange et al., 2018). However in brief, a single decision tree is formed from a series of binary splits which results in homologous or "pure" groups. The splitting process is recursive which means splitting occurs until purity is achieved if the tree is allowed to grow to its maximum depth. Decision trees make no assumptions on the input data structure (they are non-parametric), allow for interaction and collinearity among variables, and will ignore variables which are irrelevant to the dependant variable (Ziegler and König, 2013). Decision trees are fast to train, fast to make predictions, and are conceptually simple to understand. However, they suffer heavily from overfitting, an issue where the model represents the training set well, but does not generalise to sets which were not used for training (Jones and Linder, 2015). Using a model which predicts pollutant concentrations and suffers from overfitting would result in the model being contaminated with noise from the training set and unreliable predictions would impede analyses.

Random forest is an algorithm which controls for the tendency of decision trees to overfit. The algorithm achieves this by sampling (with replacement) the training set with a process called bagging (bootstrap aggregation) (Breiman, 1996). In modern usage, sampling of the independent variables is usually done during bagging too. Bagging results in a new, sampled set called out-of-bag (OOB) data. A decision tree is then grown on the OOB data. The bagging-then-tree growth is repeated, generally a few hundred times. Because OOB data is sampled, all the decision trees are grown on differing observations and independent variables which leads to a "forest" of decorrelated trees. After training, all the individual trees within the forest are used to predict, but their predictions are aggregated as a mean (or the mode for categorical dependent variables) and that forms the single ensemble prediction for the model.

The meteorological normalisation technique is pragmatic in respect to the input variables required for many common applications. Generally, routinely accessible surface meteorological variables are very effective for the process and specialised or obscure variables are generally not necessary for the technique to be applied. Although traffic counts, upper air data, and outputs from weather models will usually strengthen a model's explanatory power, the existence or access to such variables is not a prerequisite, an attribute which is very useful for most situations where such inputs are not available. For pollutants which are primarily controlled by regional scale processes, most notably particulate matter (PM) and ozone (O_3), additional variables such as boundary layer height, air mass cluster, or back trajectory information would however be beneficial to include if possible and examples can be found elsewhere, for example Grange et al. (2018).

The temporal variables used as independent variables in the meteorological normalisation models: Julian day, weekday, and hour of year are included not for their direct influence on atmospheric concentrations, but because they act as proxies for cyclical emission patterns. Hour of day for example offers a term to explain emissions with a diurnal cycle such as traffic-related rush hour emissions or domestic heating phases, while Julian day is a seasonal term which represents emissions or atmospheric chemistry which varies seasonally. These processes are generally strong drivers of concentrations of most atmospheric pollutants (Henneman et al., 2015). Random forest's ability to handle collinearity and interaction between these and the other independent variables used and the lack of need of specialised or exotic inputs results in a flexible tool kit for probing the influences of interventions on air quality time series.

1.1. Objectives

The primary objective of this paper is to apply a meteorological normalisation technique based on random forest, a machine learning algorithm to detect interventions in air quality monitoring data. This is done to gain understanding of what physical and chemical processes are driving ambient pollutant concentrations and highlight the suitability and potential of the technique to other applications.

Two case studies are presented using routine data sets in Dover, South East England where sulfur fuel limits of ships were imposed and changes in ambient sulfur dioxide (SO_2) concentrations are expected and in Central London where congestion charging and local bus fleet management has perturbed oxides of nitrogen (NO_x) emission sources. The changes in concentrations and emissions are then explained in respect to implementation of policy which would be difficult to detect with other EDA techniques where no meteorological normalisation is performed.

2. Methods

2.1. Data

2.1.1. Port of Dover SO_2

Hourly SO_2 concentrations were analysed from the Port of Dover, a major port located in Kent in the South East of England. Two air quality monitoring sites, Dover Docks and Dover Langdon Cliff's SO_2 data were queried from the Kent Air Quality database (Ricardo Energy & Environment, 2018). A nearby meteorological site, Langdon Bay located to the west of the port was used to provide surface meteorological observations and were accessed from

NOAA's Integrated Surface Database (ISD) (NOAA, 2016) (Figure 1(a)). The monitoring sites had different commissioning and decommissioning dates and neither site is still operating (Table 1). SO_2 observations are available between March 2001 and December 2012. The data capture rates for SO_2 at Dover Langdon Cliff and Dover Docks for their online period were 92 and 82 % respectively. These monitoring sites are of interest because marine fuels in British and European waters have been subject to a series of sulfur content fuel limits. The introduction and continued enforcement of these sulfur fuel limits were expected to influence ambient SO_2 concentrations. The details of these interventions are discussed further in Section 3.1.2.

Table 1: Details of the air quality monitoring sites in Dover and London used in this analysis. Sites without end dates are still operational.

Location	Site name	Site type	Latitude	Longitude	Elevation	Date start	Date end
Dover	Langdon Bay	Meteorological	51.133	1.350	117	1973-03-08	
Dover	Dover Langdon Cliff	Urban background	51.132	1.339	98	2001-03-17	2010-03-05
Dover	Dover Docks	Urban industrial	51.127	1.336	6	2006-11-17	2013-01-03
London	London Heathrow	Meteorological	51.478	-0.461	25	1948-12-01	
London	London Marylebone Road	Traffic	51.523	-0.155	35	1997-01-01	

2.1.2. London Marylebone Road NO_2 and NO_x

Hourly NO_2 and NO_x data from London's Marylebone Road air quality monitoring site were accessed from **smonitor** Europe, a European database containing the observations and metadata from the AirBase and Air Quality e-Reporting (AQER) repositories (Grange, 2016, 2017). NO_x concentrations have been monitored since July 1997 and the final year of reporting sourced from the European data repositories used was 2016. Data capture rates for NO_x and NO_2 for the analysis period were 97 %. London Heathrow, a large airport located at the far west of Greater London was used for surface meteorological observations sourced from NOAA's ISD (Figure 1(b)). London Marylebone Road is situated in a complicated central urban environment. The site is located one metre south of the kerb on the A501 trunk road and sits within an irregularly shaped street canyon. London Marylebone Road is a prominent and often analysed site due to its long observational record and diverse suite of

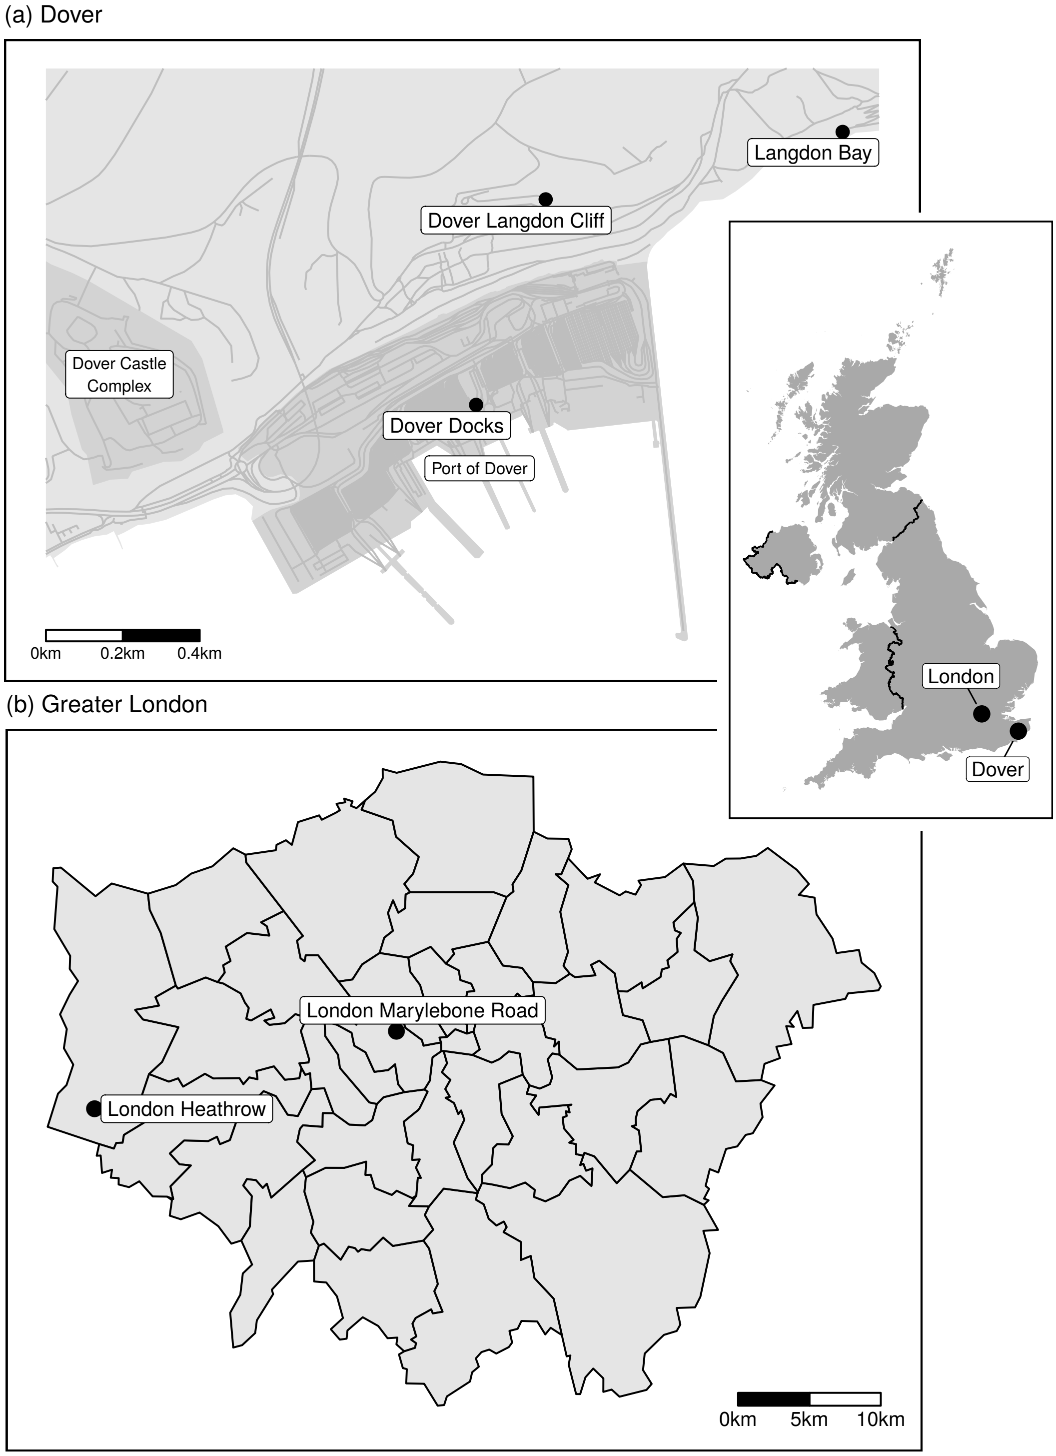

Figure 1: Maps of the study sites with a United Kingdom insert for country-scale context. The Port of Dover complex is displayed in (a) and the internal lines indicate roads and Greater London is shown in (b), with the London Boroughs and City of London indicated with internal polygons.

pollutants which are monitored at the site (Jeanjean et al., 2017).

NO_x and NO_2 concentrations across European cities are a significant issue and many member states are non-compliant to the legal European ambient air quality limits (Weiss et al., 2012; Grange et al., 2017). Almost all locations which are non-compliant are classified as roadside (or 'traffic-influenced') (European Environment Agency, 2016). London has some of the highest roadside concentrations of NO_x and NO_2 in Europe and London Marylebone Road (Figure 1(b)) is an often referenced monitoring site for its high concentrations.

To combat the issue of traffic congestion, Greater London authorities imposed the Congestion Charge Zone (CCZ), which was first enforced in February 2003 (Atkinson et al., 2009). Since that time, the London Low Emission Zone (LEZ), and the Emissions Surcharge (better known as the T-Charge) have also been implemented to combat air pollution (Transport for London, 2018). The details and start dates of these various measures are displayed in Table 2. All these interventions are significant investments with large amounts of planning and resources to execute and maintain.

Table 2: Details of interventions within Greater London to counter traffic congestion.

Name	Abbreviation	Start date	Area covered	Operation
Congestion Charge Zone	CCZ	2003-02-17	Central London	07:00–18:00 Mo-Fr
London Low Emission Zone (first phase)	LEZ	2008-02-04	Greater London	24/7
London Low Emission Zone (second phase)	LEZ	2012-01-03	Greater London	24/7
Emissions Surcharge	T-Charge	2017-10-23	Central London	07:00–18:00 Mo-Fr
Ultra Low Emission Zone (planned)	ULEZ	2019-04-08	Central London	24/7

2.2. Modelling and the hyperparameters

For both examples, the meteorological normalisation procedure was conducted in the same way and the **rmweather** R package (version 0.1.2) was used for this process (R Core Team, 2018; Grange, 2018). The number of trees for the random forest models was fixed at 300, the minimal node size was five, and the number of variables split at each node was the default for regression mode: the rounded down square root of the number of independent variables which in these examples was three (**rmweather**'s function arguments `n_trees`, min_node_size, and mtry respectively). The independent variables used were: Unix date (number of seconds since 1970-01-01) as the trend term, Julian day as the seasonal term, weekday, hour of day, air temperature, relative humidity, wind direction, wind speed, and atmospheric pressure. Training was only conducted on observations which had non-missing wind speed and the pollutant being modelled. Three hundred predictions were used to calculate the meteorologically normalised trend. The normalised trends were aggregated to monthly resolution for presentation in Section 3. A conceptual representation of the meteorological normalisation processes is displayed in Figure A1.

For the Dover SO_2 examples, models were calculated using the full observational set, but after investigating the models (discussed in Section 3.1.1), the observations were filtered to wind directions which were sourced from the port and these models are the ones which were used for the time series analysis (Section 3.1.2). For observations at London Marylebone Road, no filtering was undertaken. In the case of London Marylebone Road, there are a large number of potential events which could influence pollutant concentrations and emissions. To objectively identify events, the meteorologically normalised time series were tested for breakpoints or changes in structure. The structural change algorithm is described in Zeileis et al. (2002); Zeileis et al. (2003) and was implemented with the **strucchange** R package.

The random forest algorithm does not directly offer the ability to determine error or uncertainty of estimates. However, uncertainty is important to consider in many situations. To enable uncertainty to be evaluated for the case studies, 50 random forest models were grown for each example with the hyperparameters described above, but with randomly sampled (bootstrapped) input sets. The bootstrapping of the observational data ensured the models were grown on different training sets. The importance values (a measure of the variables' strength or influence on prediction), partial dependencies, and predictions for each of the 50 models were then summarised. The summaries used from the "ensemble of the ensembles" were the mean, and the 2.5 % and 97.5 % quantiles of the 50 estimates *i.e.* a range that spans the 95 % confidence interval in the mean. The model performance statistics for the four sets of models are displayed in Table 3.

Table 3: Mean random forest model performance statistics four the four sets of models grown for the analysis.

Location	Model	n	R^2
Dover	Dover Docks SO_2	34224	0.67
Dover	Dover Langdon Cliff SO_2	53535	0.63
London	London Marylebone Road NO_2	131677	0.82
London	London Marylebone Road NO_x	131677	0.83

3. Results and discussion

3.1. Port of Dover SO_2

3.1.1. Models

The random forest models grown for SO_2 at the two Dover sites had R^2 values of 63 and 67 % (Table 3), therefore, the models had moderate explanatory ability for Dover's SO_2 concentrations. However, it should be noted that predicting concentrations over such short time periods with intermittent source strength is challenging and data capture was less than ideal for these monitoring sites. The moderate performance can be explained by SO_2 at this location containing large amounts of variation due to ship movements and if winds were in a favourable direction to transport emissions from the port complex to the monitoring sites (southerlies). Indeed, wind direction was the most important variable for SO_2 explanation for the random forest models (Figure 2).

Partial dependence plots of decision tree models allow the learning process to be interpreted and a data user to examine how variables are being handled in the predictive model. Figure 3 demonstrates a two-way partial dependence plot for SO_2 concentrations at Dover Landon Cliff using wind direction and date (the trend term) as the independent variables. The feature which is most clear is the band of increased SO_2 dependence between 150 and 210 degrees. Outside of this band of southerly winds, there were low levels of dependence on SO_2 concentrations. The Dover Landon Cliff monitoring site was located north of the Port of Dover docks and very slightly to the east (Figure 1(a)). The partial dependence on wind direction is consistent with this location and indicates that wind direction was

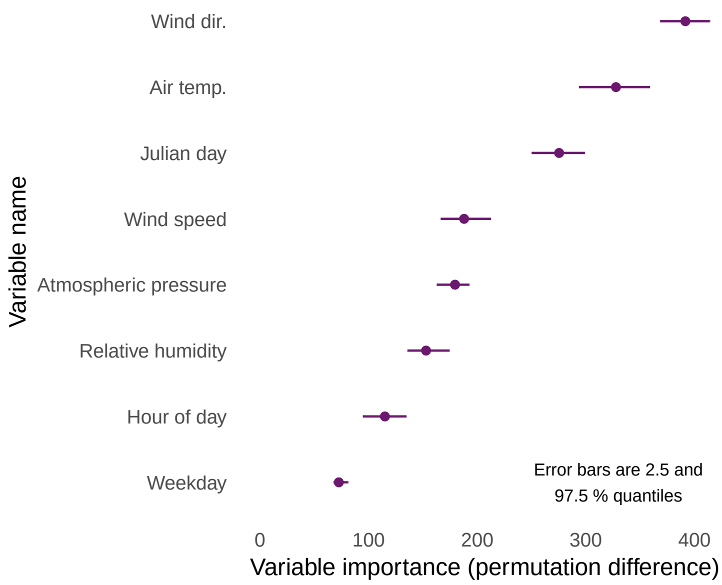

Figure 2: Variable importance plot for SO_2 at Dover Langdon Cliff between 2001 and 2010 calculated by 50 random forest models.

handled sensibly in the random forest model. This observation can be confirmed further with a bivariate polar plot of mean SO_2 concentrations by wind direction and speed at the monitoring site (Figure 4). The first sulfur content fuel change in mid-August 2006 can also be seen in the two-way partial dependence plot as a clear reduction in SO_2 dependence when winds were sourced from the port (the south; discussed further in Section 3.1.2; Figure 3).

Another clear feature isolated by the partial dependence plots was that SO_2 concentrations increased with increasing air temperature at the Dover monitoring sites (Figure 5). This relationship was an unexpected outcome because generally, pollutant concentrations are inversely related to air temperature because emissions are more efficiently diluted during warmer periods owing to increased thermal turbulence. For some sources such as heating, emissions are greater at lower temperatures, but when considering shipping emissions, this would be negligible. At Dover, the SO_2 relationship between concentrations and air temperatures was indicative of convective thermal mixing being an important physical process which resulted in SO_2 emitted by ships to be mixed towards the measurement site at the cliff top. This turbulent mixing at high temperatures resulted in high SO_2 concentrations at

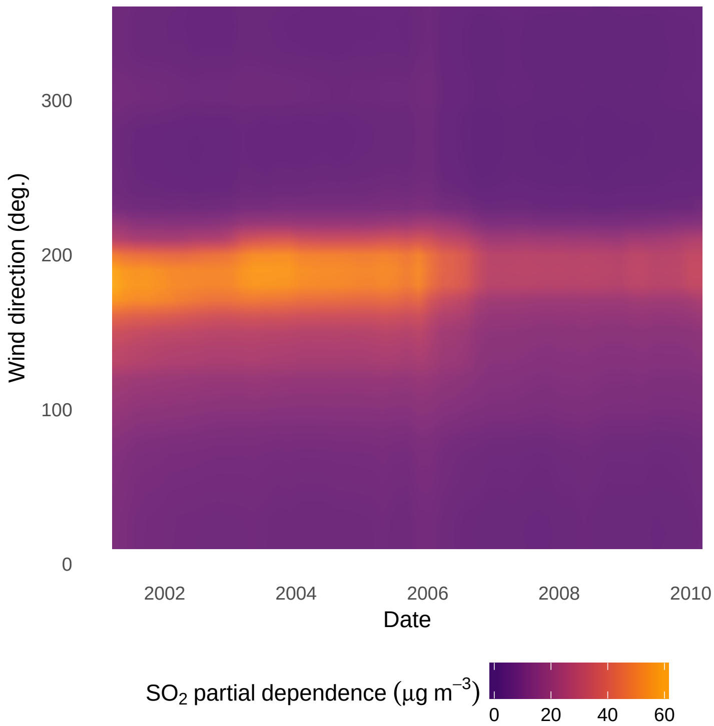

Figure 3: Partial dependence of wind direction and date on SO_2 concentrations at Dover Landon Cliff between 2001 and 2010. The Dover Landon Cliff monitoring site was located north of the Port of Dover (Figure 1(a)).

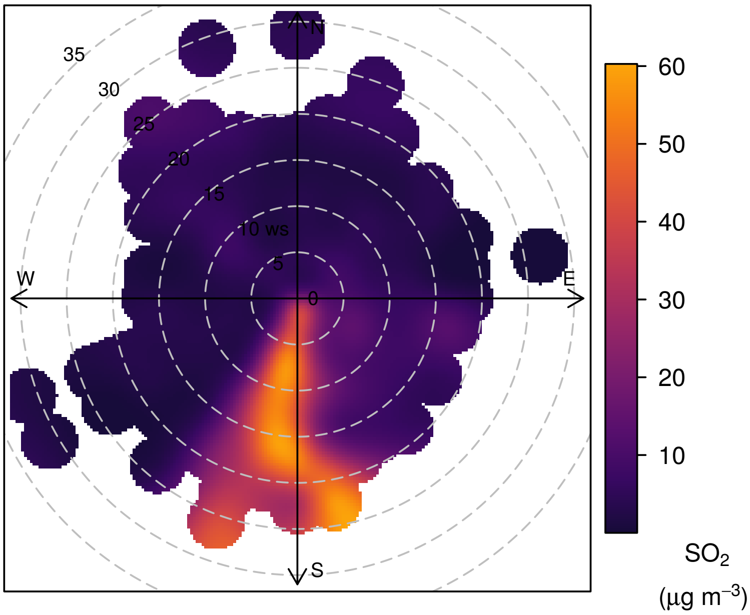

Figure 4: Bivariate polar plot of mean hourly SO_2 concentrations at Dover Landon Cliff between 2001 and 2010. The Dover Landon Cliff monitoring site was located north of the Port of Dover (for a location map, see Figure 1(a)).

the surface and this feature cannot be easily observed in the hourly observational data. The illumination of such physical processes is a major advantage of the random forest algorithm compared to other machine learning methods such as support vector machines (SVM) or artificial neural networks (ANNs) because they do not offer the same amount of model legibility.

3.1.2. Influence of sulfur fuel limits on SO_2 concentrations

Since the early 2000s, there has been a number of increasingly stringent sulfur based fuel limits imposed on ships operating in British and European Union (EU) waters due to their status as Sulfur Emission Control Areas (SECAs) or Emission Control Areas (ECAs). The most important events for sulfur control were implemented on August 11, 2006 and January 1, 2010. In August 2006, the MARPOL Annex IV regulations were applied which introduced a 1.5 % sulfur limit on fuel oils used by vessels moving between EU ports (International Maritime Organization, 2005). The pre-August 2006 sulfur content for British vessels has been estimated at 2.7 % which represents a reduction in sulfur content of 44 % (Entec, 2010).

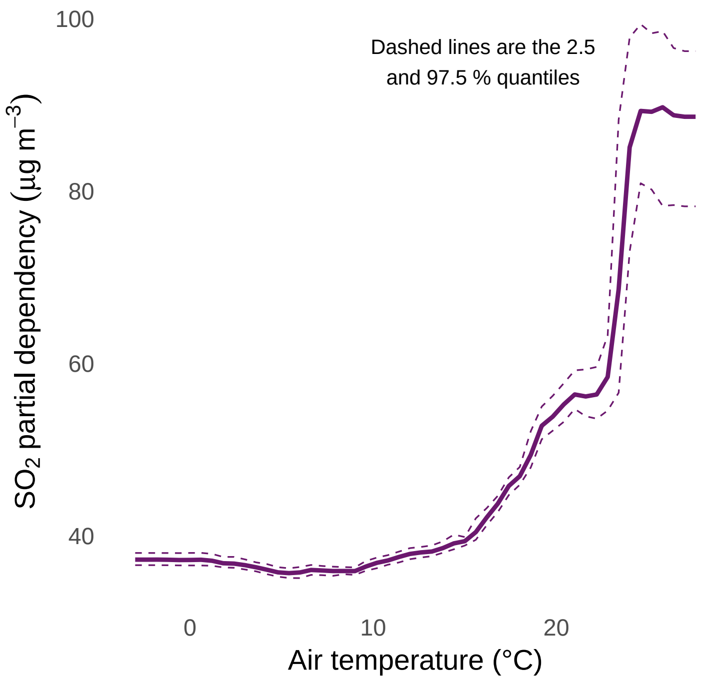

Figure 5: Partial dependence of SO_2 on air temperature at Dover Landon Cliff between 2001 and 2010 calculated by 50 random forest models.

At the start of 2010 an additional limit was imposed for all vessels at berth where such vessels were required to be operated with maximum fuel sulfur content of 1 %. These changes should be evident in the SO_2 time series of the nearby ambient monitoring sites. However, if a time series is plotted, the influence of these changes are subtle and not clear due to the high amounts of variation within SO_2 concentrations (Figure 6).

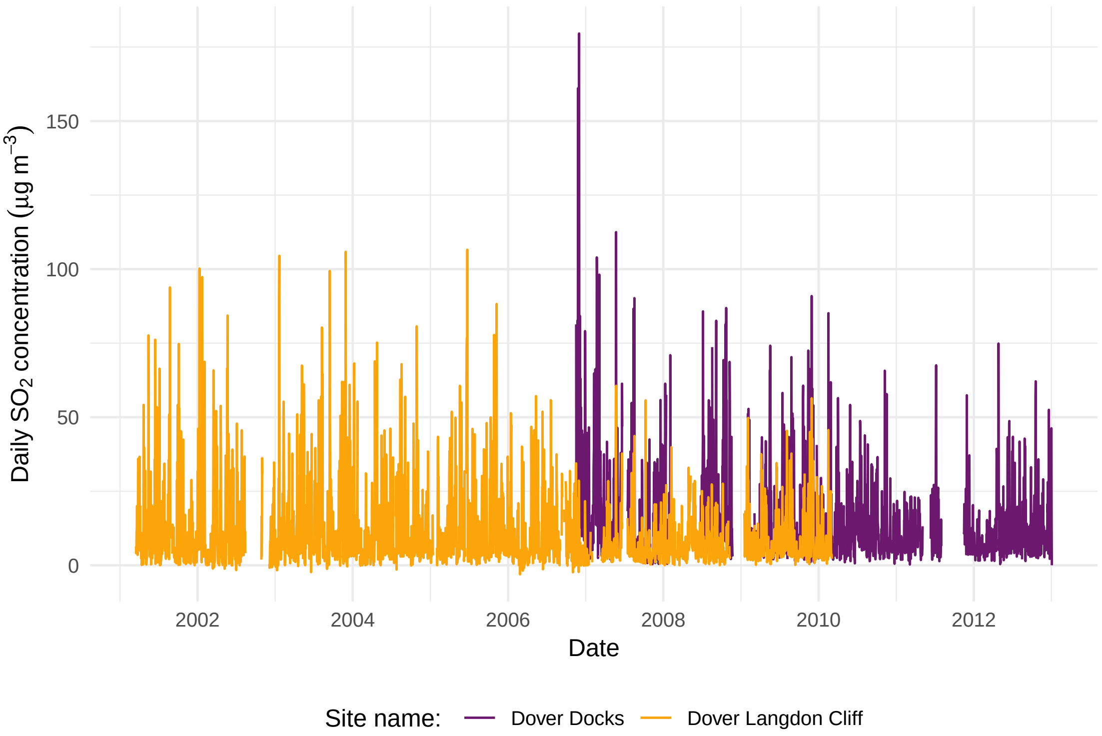

Figure 6: Daily SO_2 concentrations at two monitoring sites in Dover between 2001 and 2012.

The meteorologically normalised SO_2 time series for the Dover sites are displayed in Figure 7, after the observations were filtered to wind directions which came for the port, hence the tight 95 % confidence intervals. The dates when changes in sulfur fuel content were implemented are displayed as vertical lines in Figure 7 and the influence of sulfur fuel changes are clear (compared with Figure 6).

At Dover Langdon Cliff, the monitoring site which was online during the MARPOL 1.5 % fuel sulfur limit transition during August 2001 shows the shift in ambient SO_2 very clearly (Figure 7). The mean meteorologically normalised SO_2 concentrations for the pre- and

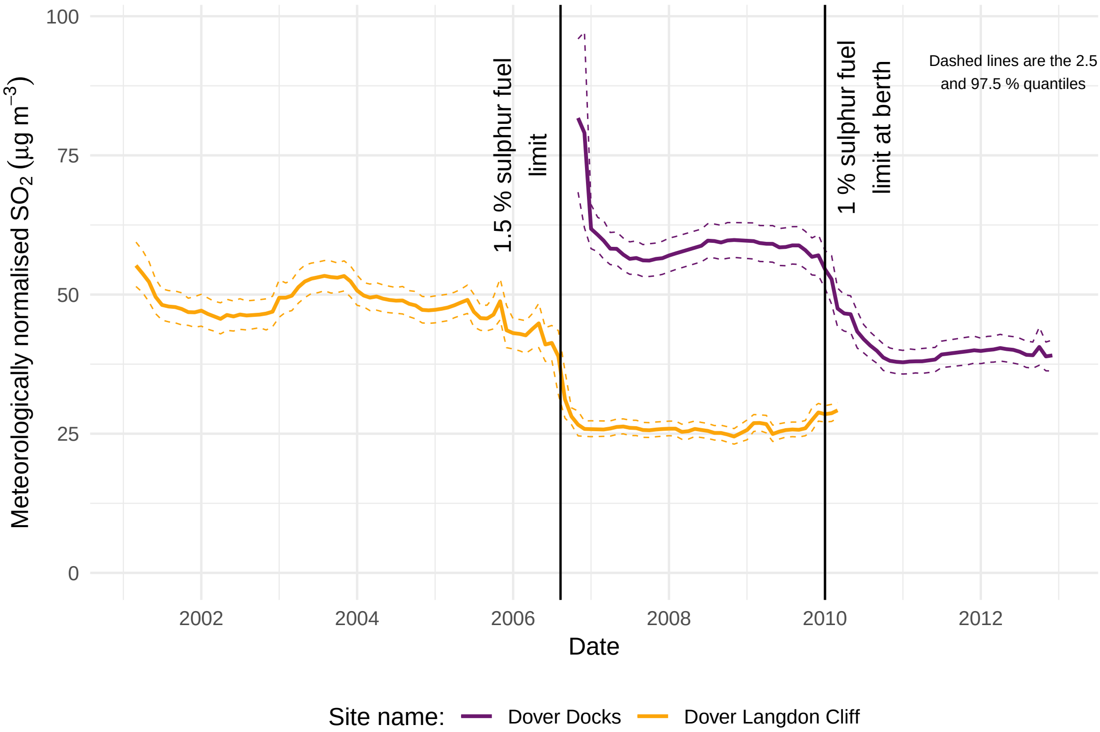

Figure 7: Meteorologically normalised SO_2 concentrations at two monitoring sites in Dover between 2001 and 2012 as calculated by 50 random forest models. The vertical lines show the start dates of when changes in marine sulfur fuel content were implemented.

post-fuel change periods were 48 and 26 $\mu g\, m^{-3}$ respectively. This difference represented in percentage change is 45 % and the corresponding estimated change in sulfur fuel content was 44 %. This extremely good agreement between sulfur content fuel changes and normalised ambient SO_2 concentrations suggests that the Port of Dover activities and ship movements remained constant during the transition phase and the source of SO_2 at this location was almost exclusively from the port.

The second sulfur fuel content change was implemented on January 1, 2010 and this intervention is also clearly displayed in the meteorologically normalised SO_2 concentrations of the Dover Docks monitoring site (Figure 7). The percentage change in fuel sulfur content was 33 % and the percentage change in ambient SO_2 concentrations was 32 %. Like the previous intervention, these two percentage changes match almost exactly, which is somewhat surprising because the intervention was applied only to berthed vessels which would only make up a component of the Port of Dover activities.

3.2. London Marylebone Road NO_x

3.2.1. Models

The random forest models of NO_x and NO_2 at London Marylebone Road performed well and had R^2 values of 82 and 83 % respectively (Table 3). This good performance can be explained by hour of day being a very good predictor for traffic flows and therefore emissions at this location for these (mostly) traffic-sourced pollutants (Figure 8). The performance of the random forest models would be rather difficult to achieve with dispersion or deterministic models in such a complicated location. For example, the dispersion models evaluated in Carslaw et al. (2013) struggled to represent the street canyon environment, even when traffic information was taken into account. The importance plots for the London Marylebone Road models also show that wind direction is the most important variable to predict NO_2 and NO_x concentrations. London Marylebone Road is located in a street canyon and is subjected to complex flows, including ventilation, vortices, and leeward accumulation of pollutants, (primarily) dependent on wind direction (Carslaw and Carslaw, 2007; Catalano et al., 2016). This complexity is demonstrated in the importance of wind direction in explaining NO_x and

NO_2 concentrations (Figure 8) and this has been noted before at this location (Charron and Harrison, 2005; Westmoreland et al., 2007).

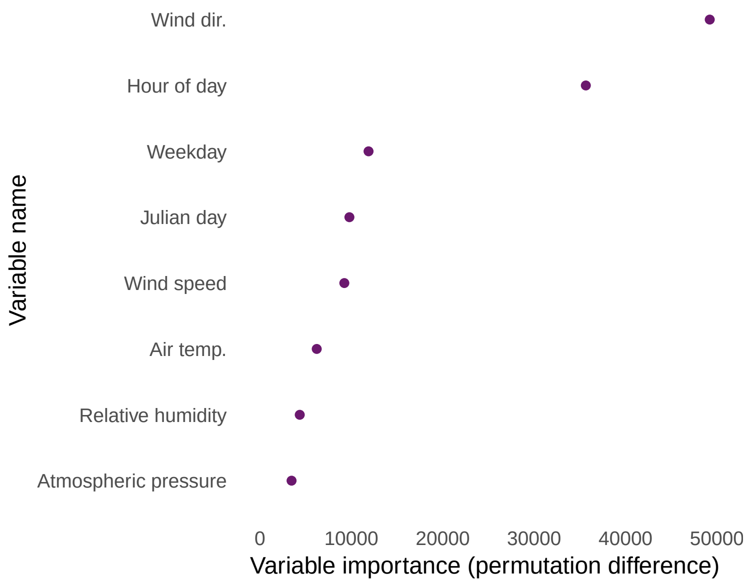

Figure 8: Variable importance plot for 50 NO_2 random forest models for London Marylebone Road. The uncertainty among the importances of the 50 models was very small and therefore the quantiles are not shown. The importances for the NO_x models were very similar.

3.2.2. Changes in primary NO_2

Using the predictive models for meteorological normalisation results in very clear and almost noiseless meteorologically normalised trends shown in Figure 9. It is immediately clear that NO_x and NO_2 are not behaving the same way at this monitoring location. This is because of changes in vehicular primary (directly emitted) NO_2 during the analysis period (1997–2016) (Carslaw, 2005; Carslaw et al., 2016; Grange et al., 2017). The vertical lines on Figure 9 show the breakpoints identified by structural change analysis after the meteorological normalisation procedure.

NO_x concentrations decreased after the introduction of a bus lane adjacent to the monitoring site in 2001 but have remained near constant since the introduction of the CCZ in February 2003 (Figure 9 and Table 2). Despite the progressively stringent vehicular emission controls being applied across Europe between 2003 and 2016 (the last year of data in analysis), they have had little effect to NO_x at London Marylebone Road. This observation could be, at least partly, explained by the disconnect between laboratory testing and real-world emissions of NO_x which become a public issue after the diesel emission scandal in September 2015 (Brand, 2016; Schmidt, 2016). However, heavy duty vehicles are also very important to consider alongside passenger vehicles at this Central London location (Laybourn-Langton et al., 2016; Greater London Authority, 2017).

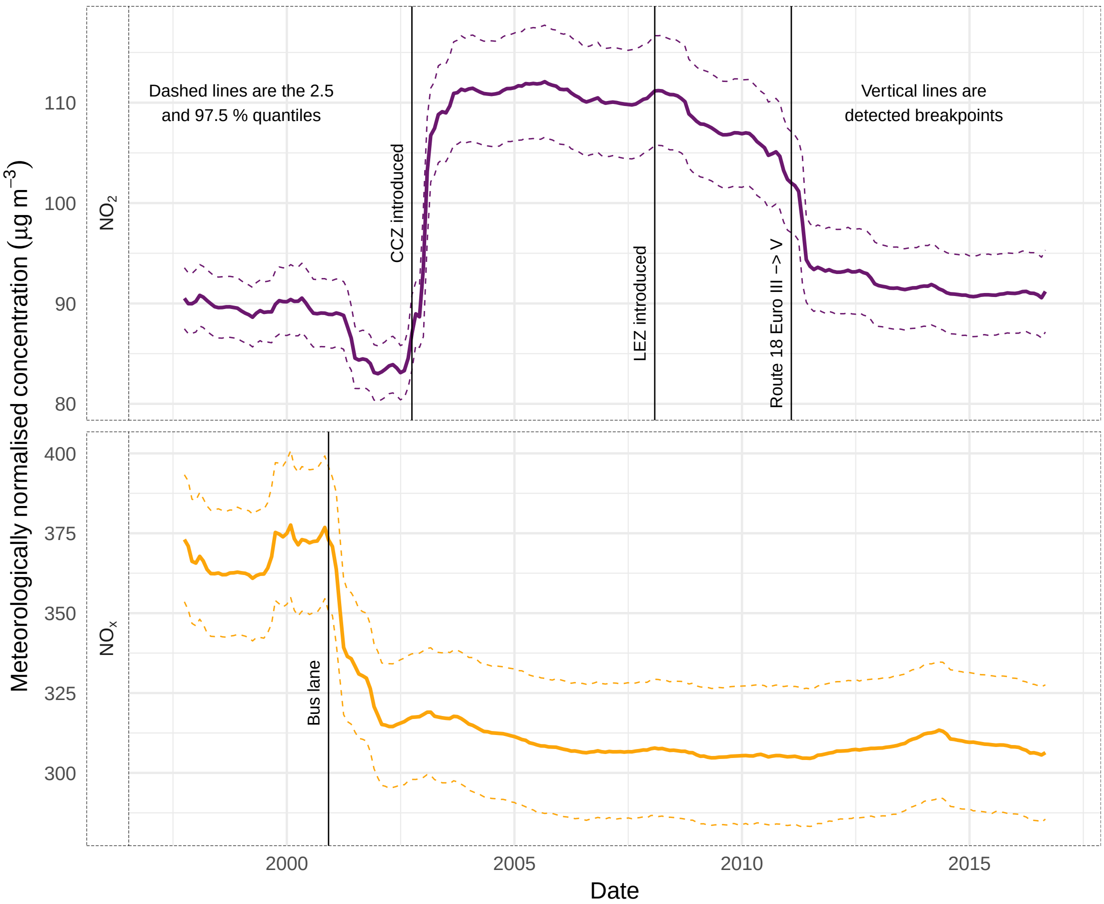

Figure 9: Meteorologically normalised NO_x and NO_2 at London Marylebone Road between 1997 and 2016 as calculated by 50 random forest models (for each pollutant). The vertical lines on show the breakpoints identified by structural change analysis.

NO_2 concentrations at London Marylebone Road have increased since 1997 and were at their maximum between 2002 and 2008 (Figure 9). The changes observed can be explained by changes to the vehicle fleet using the adjacent A501 road resulting from the introduction of congestion charging, London's Low Emission Zone, and evolution of the local bus fleet. The rapid increase of NO_2 concentrations was observed in the meteorologically normalised time series between July 2002 and July 2003 (Figure 9). The CCZ was introduced in mid-February 2002; right in the middle of the period of increasing NO_2 and within six months of the suggested breakpoint (October 2012). The increase in NO_2 concentrations was due to increased primary NO_2 because no change in the meteorologically normalised NO_x was observed at the same time.

The implementation of the CCZ was accompanied with a retrofitting programme of Euro III local buses with continuously regenerating diesel particulate filters (CRDPF, also known by their commercial name: CRT filters). CRDPF are passive devices and have two components: an upstream oxidation catalyst and a particulate matter (PM) filter. The oxidation catalyst oxidises NO within the exhaust stream to NO_2 and this NO_2 is then used as a PM oxidant in the filter-proper. The observations show that these retrofitted passive devices were not optimised because much of the generated NO_2 was not reduced within the PM filter and was therefore emitted into the roadside atmosphere and thus significantly increased ambient NO_2 concentrations (Figure 9).

NO_2 concentrations remained approximately stable until February 2008 when London's Low Emission Zone (LEZ) was introduced and NO_2 concentrations began to decrease (Figure 9). The second NO_2 breakpoint was detected for February 2008 giving some evidence that the LEZ reduced NO_2 concentrations at London Marylebone Road (although no corresponding change in NO_x was observed). However, during this period the local bus fleets were also being progressively replaced with newer buses compliant to the later Euro IV, V, and VI heavy vehicle emission standards (Finn Coyle, Tom Cunnington, and Gabrielle Bowden (Transport for London), personal communication, March 2018) as well of natural vehicle turnover removing older and more polluting vehicles from the in-service fleet. The third NO_2 breakpoint identified coincided with route 18, the bus route with the highest peak vehicle requirements (PVR), shifting from Euro III to Euro V vehicles in late 2010 (Figure 9). After 2011, NO_2 concentrations continued to decline with the introduction of Euro VI and hybrid buses servicing the 453, 27, and 205 routes. By the end of 2016, NO_2 had declined to almost pre-CCZ concentrations. The features displayed in the normalised time series were not clear in the raw concentration data (displayed in Figure A2) and the breakpoints identified were unable to be resolved without the meteorological normalisation technique.

The tandem use of the meteorological normalisation procedure and breakpoint analysis is powerful and can revel many changes, but in many cases there may not be sufficient information or metadata to help explain the changes observed. In this Central London example, many of the factors driving pollutant concentrations are known due to the site's prominence.

London Marylebone Road also monitors ozone (O_3), something which is rare for roadside monitoring locations in Europe. The NO_2, NO_x, and O_3 complement allows for the estimation of primary NO_2 with an independent method by determining the total oxidant (OX; NO_2 + O_3) within NO_x (Jenkin, 2004; Carslaw and Beevers, 2005). Figure 10 shows monthly estimates of the primary NO_2 fraction at London Marylebone Road with robust linear regression. Figure 10 is consistent with Figure 9 with a rapid increase in primary NO_2 during 2002 and a reduction, but at a slower rate after 2008 thus further confirming and validating that the trends observed in Figure 9 are driven by changes in primary NO_2 emissions. The reason why the trend is similar in Figure 10 and Figure 9 is that at this particular site increased emissions of primary NO_2 were sufficient to have a measurable effect on ambient concentrations.

4. Conclusions

Controlling for changes of meteorology is an important component to consider when conducting air quality data analysis over time. A meteorological normalisation technique using random forest was used to investigate interventions in routine air quality monitoring data from two areas. The interventions applied to marine fuel content changes were explored in Dover, a port city in the South East of England and the interventions were represented in the meteorologically normalised time series almost exactly. The non-black box nature of the random forest models was used to investigate the dependence of pollutant concentrations

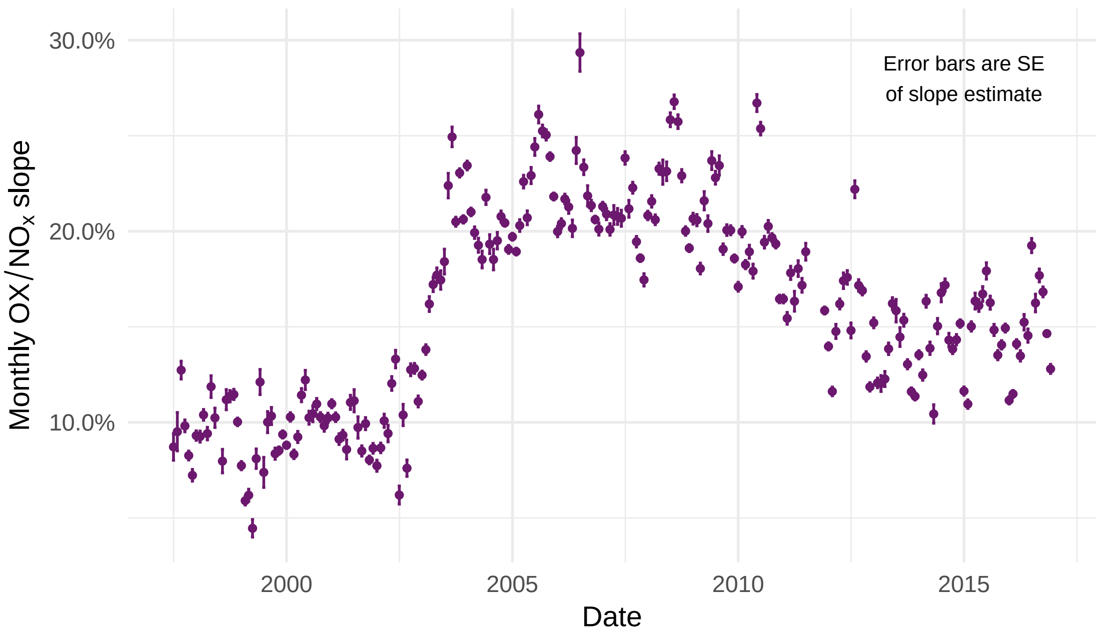

Figure 10: Monthly total oxidant (OX; NO_2 + O_3) at London Marylebone Road between 1997 and 2016. Slope and errors were calculated with robust linear regression.

on meteorological variables such as air temperature and wind direction which highlighted the benefit of the technique where physical and chemical atmospheric processes can be illuminated, understood, and explained.

In the example of the implementation of congestion charging in Central London, very clear changes in primary NO_2 emissions were displayed in the meteorologically normalised time series. The performance of these roadside models was high due to the models' ability to use wind direction and hour of day very effectively, something which dispersion or deterministic models struggle with when used for modelling street canyon environments. The case studies presented are both examples where there is significant ability to cross check the observed features with available information on changes in the sites' local environments to validate the outputs.

The meteorological normalisation technique is very relevant for exploring the influence of interventions or management activities on local air quality. The combination of a non-parametric method, the lack of need for specialised measurements, and the effective use of proxy variables lends the technique to a wide range of air quality data analysis applications.

Acknowledgements

S.K.G. was supported by Anthony Wild with the provision of the Wild Fund Scholarship. This work was also partially funded by Natural Environment Research Council (NERC) [grant number: NE/N007115/1].

Competing interests

The authors declare no competing interest.

Highlights

- Detecting the influence of air quality interventions is important
- Changes in meteorology over time complicates air quality intervention analysis
- Meteorological normalisation was applied in two locations to explore interventions
- The changes detected in the normalised time series were associated to interventions
- The non-black-box nature of the procedure allows for interpretation of results

References

Anh, V., Duc, H., Azzi, M., 1997. Modeling anthropogenic trends in air quality data. Journal of the Air & Waste Management Association 47 (1), 66–71.
URL https://doi.org/10.1080/10473289.1997.10464406

Atkinson, R., Barratt, B., Armstrong, B., Anderson, H., Beevers, S., Mudway, I., Green, D., Derwent, R., Wilkinson, P., Tonne, C., Kelly, F., Nov. 2009. The impact of the congestion charging scheme on ambient air pollution concentrations in London. Atmospheric Environment 43 (34), 5493–5500.
URL http://www.sciencedirect.com/science/article/pii/S1352231009006268

Barmpadimos, I., Hueglin, C., Keller, J., Henne, S., Prévôt, A. S. H., Feb. 2011. Influence of meteorology on PM_{10} trends and variability in Switzerland from 1991 to 2008. Atmospheric Chemistry and Physics 11 (4), 1813–1835.
URL http://www.atmos-chem-phys.net/11/1813/2011/

Brand, C., 2016. Beyond 'Dieselgate': Implications of unaccounted and future air pollutant emissions and energy use for cars in the United Kingdom. Energy Policy 97, 1–12.
URL http://www.sciencedirect.com/science/article/pii/S030142151630341X

Breiman, L., 1996. Bagging predictors. Machine Learning 24 (2), 123–140.
URL http://dx.doi.org/10.1007/BF00058655

Breiman, L., 2001. Random forests. Machine Learning 45 (1), 5–32.
URL http://dx.doi.org/10.1023/A:1010933404324

Carslaw, D., Apsimon, H., Beevers, S., Brookes, D., Carruthers, D., Cooke, S., Kitwiroon, N., Oxley, T., Stedman, J., Stocker, J., 2013. Defra Phase 2 urban model evaluation, UK AIR: Air Information Resource.
URL https://uk-air.defra.gov.uk/library/reports?report_id=777

Carslaw, D. C., Aug. 2005. Evidence of an increasing NO_2/NO_x emissions ratio from road traffic emissions. Atmospheric Environment 39 (26), 4793–4802.
URL http://www.sciencedirect.com/science/article/pii/S1352231005005443

Carslaw, D. C., Beevers, S. D., Jan. 2005. Estimations of road vehicle primary NO_2 exhaust emission fractions using monitoring data in London. Atmospheric Environment 39 (1), 167–177.
URL http://www.sciencedirect.com/science/article/pii/S1352231004008775

Carslaw, D. C., Carslaw, N., 2007. Detecting and characterising small changes in urban nitrogen dioxide concentrations. Atmospheric Environment 41 (22), 4723–4733.
URL http://www.sciencedirect.com/science/article/pii/S1352231007002919

Carslaw, D. C., Murrells, T. P., Andersson, J., Keenan, M., 2016. Have vehicle emissions of primary NO_2 peaked? Faraday Discussions 189 (0), 439–454.
URL http://dx.doi.org/10.1039/C5FD00162E

Carslaw, D. C., Ropkins, K., Bell, M. C., Nov. 2006. Change-point detection of gaseous and particulate traffic-related pollutants at a roadside location. Environmental Science & Technology 40 (22), 6912–6918.
URL http://dx.doi.org/10.1021/es060543u

Catalano, M., Galatioto, F., Bell, M., Namdeo, A., Bergantino, A. S., Jun. 2016. Improving the prediction of air pollution peak episodes generated by urban transport networks. Environmental Science & Policy 60, 69–83.
URL http://www.sciencedirect.com/science/article/pii/S1462901116300594

Charron, A., Harrison, R. M., Oct. 2005. Fine ($PM_{2.5}$) and Coarse ($PM_{2.5-10}$) Particulate Matter on A Heavily Trafficked London Highway: Sources and Processes. Environmental Science & Technology 39 (20), 7768–7776.
URL http://dx.doi.org/10.1021/es050462i

Derwent, R., Middleton, D., Field, R., Goldstone, M., Lester, J., Perry, R., 1995. Analysis and interpretation

of air quality data from an urban roadside location in Central London over the period from July 1991 to July 1992. Atmospheric Environment 29 (8), 923 – 946.
URL http://www.sciencedirect.com/science/article/pii/135223109400219B

Entec, 2010. Defra UK Ship Emissions Inventory—Final Report, doc Reg No. 21897-01.
URL https://uk-air.defra.gov.uk/assets/documents/reports/cat15/1012131459_21897_Final_Report_291110.pdf

European Environment Agency, 2016. Air quality in Europe — 2016 report. EEA Report. No 28/2016.
URL http://www.eea.europa.eu/publications/air-quality-in-europe-2016

Friedman, J., Hastie, T., Tibshirani, R., 2001. The Elements of Statistical Learning. Data Mining, Inference, and Prediction, 2nd Edition. Vol. 1. Springer series in statistics Springer, Berlin.

Grange, S. K., 2016. **smonitor**: A framework and a collection of functions to allow for maintenance of air quality monitoring data.
URL https://github.com/skgrange/smonitor

Grange, S. K., 2017. Technical note: **smonitor** Europe. Tech. rep., Wolfson Atmospheric Chemistry Laboratories, University of York.
URL https://doi.org/10.13140/RG.2.2.20555.49448/1

Grange, S. K., 2018. **rmweather**: Tools to Conduct Meteorological Normalisation on Air Quality Data. R package version 0.1.2.
URL https://CRAN.R-project.org/package=rmweather

Grange, S. K., Carslaw, D. C., Lewis, A. C., Boleti, E., Hueglin, C., May 2018. Random forest meteorological normalisation models for Swiss PM_{10} trend analysis. Atmospheric Chemistry and Physics 18 (9), 6223–6239.
URL https://www.atmos-chem-phys.net/18/6223/2018/

Grange, S. K., Lewis, A. C., Moller, S. J., Carslaw, D. C., Dec. 2017. Lower vehicular primary emissions of NO_2 in Europe than assumed in policy projections. Nature Geoscience 10 (12), 914–918.
URL https://doi.org/10.1038/s41561-017-0009-0

Greater London Authority, 2017. London atmospheric emissions inventory (laei) 2013.
URL https://data.london.gov.uk/dataset/london-atmospheric-emissions-inventory-2013

Henneman, L. R., Holmes, H. A., Mulholland, J. A., Russell, A. G., Oct. 2015. Meteorological detrending of primary and secondary pollutant concentrations: Method application and evaluation using long-term (2000–2012) data in Atlanta. Atmospheric Environment 119, 201–210.
URL http://www.sciencedirect.com/science/article/pii/S1352231015302521

International Maritime Organization, 2005. Revised MARPOL Annex VI, annex VI of MARPOL addresses air pollution from ocean-going ships.

URL http://www.imo.org/en/OurWork/Environment/PollutionPrevention/AirPollution/Pages/Air-Pollution.aspx

Jeanjean, A. P. R., Buccolieri, R., Eddy, J., Monks, P. S., Leigh, R. J., Mar. 2017. Air quality affected by trees in real street canyons: The case of Marylebone neighbourhood in central London. Urban Forestry & Urban Greening 22, 41–53.
URL http://www.sciencedirect.com/science/article/pii/S1618866716303740

Jenkin, M. E., Sep. 2004. Analysis of sources and partitioning of oxidant in the UK—Part 2: contributions of nitrogen dioxide emissions and background ozone at a kerbside location in London. Atmospheric Environment 38 (30), 5131–5138.
URL http://www.sciencedirect.com/science/article/pii/S1352231004006193

Jones, Z., Linder, F., 2015. Exploratory data analysis using random forests, 73rd annual MPSA conference, April 16-19, 2015, Chicago, United States of America.
URL https://pdfs.semanticscholar.org/e7b7/3565b07a7f1369a20b1055f222423f0feb34.pdf

Laybourn-Langton, L., Quilter-Pinner, H., Ho, H., 2016. Lethal and illegal: Solving london's air pollution crisis, institute for Public Policy Research.
URL http://www.ippr.org/read/lethal-and-illegal-solving-londons-air-pollution-crisis

Libiseller, C., Grimvall, A., Waldén, J., Saari, H., 2005. Meteorological normalisation and non-parametric smoothing for quality assessment and trend analysis of tropospheric ozone data. Environmental Monitoring and Assessment 100 (1), 33–52.
URL http://dx.doi.org/10.1007/s10661-005-7059-2

Lyubchich, V., Gel, Y. R., El Shaarawi, A., 2013. On detecting non-monotonic trends in environmental time series: a fusion of local regression and bootstrap. Environmetrics 24 (4), 209–226.
URL http://dx.doi.org/10.1002/env.2212

Monks, P., Granier, C., Fuzzi, S., Stohl, A., Williams, M., Akimoto, H., Amann, M., Baklanov, A., Baltensperger, U., Bey, I., Blake, N., Blake, R., Carslaw, K., Cooper, O., Dentener, F., Fowler, D., Fragkou, E., Frost, G., Generoso, S., Ginoux, P., Grewe, V., Guenther, A., Hansson, H., Henne, S., Hjorth, J., Hofzumahaus, A., Huntrieser, H., Isaksen, I., Jenkin, M., Kaiser, J., Kanakidou, M., Klimont, Z., Kulmala, M., Laj, P., Lawrence, M., Lee, J., Liousse, C., Maione, M., McFiggans, G., Metzger, A., Mieville, A., Moussiopoulos, N., Orlando, J., O'Dowd, C., Palmer, P., Parrish, D., Petzold, A., Platt, U., Pschl, U., Prvt, A., Reeves, C., Reimann, S., Rudich, Y., Sellegri, K., Steinbrecher, R., Simpson, D., ten Brink, H., Theloke, J., van der Werf, G., Vautard, R., Vestreng, V., Ch. Vlachokostas, von Glasow, R., 2009. Atmospheric composition change - global and regional air quality. Atmospheric Environment 43 (33), 5268 – 5350.
URL http://www.sciencedirect.com/science/article/B6VH3-4X3N46N-1/2/

1db0fa3c5afafc9418ab227802a71755
NOAA, 2016. Integrated Surface Database (ISD).
URL https://www.ncdc.noaa.gov/isd
R Core Team, 2018. R: A Language and Environment for Statistical Computing. R Foundation for Statistical Computing, Vienna, Austria.
URL https://www.R-project.org/
Ricardo Energy & Environment, 2018. Kent air quality database.
URL http://www.kentair.org.uk
Schmidt, C. W., Jan. 2016. Beyond a One-Time Scandal: Europe's Onging Diesel Pollution Problem. Environmental Health Perspectives 124 (1), A19–A22.
URL http://www.ncbi.nlm.nih.gov/pmc/articles/PMC4710587/
Stull, R. B., 1988. An Introduction to Boundary Layer Meteorology. Kluwer Academic Publishers, London.
Tong, W., Hong, H., Fang, H., Xie, Q., Perkins, R., Mar. 2003. Decision forest: Combining the predictions of multiple independent decision tree models. Journal of Chemical Information and Computer Sciences 43 (2), 525–531.
URL http://dx.doi.org/10.1021/ci020058s
Transport for London, 2018. Driving.
URL https://tfl.gov.uk/modes/driving/
Weiss, M., Bonnel, P., Kühlwein, J., Provenza, A., Lambrecht, U., Alessandrini, S., Carriero, M., Colombo, R., Forni, F., Lanappe, G., Le Lijour, P., Manfredi, U., Montigny, F., Sculati, M., Dec. 2012. Will Euro 6 reduce the NO_x emissions of new diesel cars? — Insights from on-road tests with Portable Emissions Measurement Systems (PEMS). Atmospheric Environment 62, 657–665.
URL http://www.sciencedirect.com/science/article/pii/S1352231012008412
Westmoreland, E. J., Carslaw, N., Carslaw, D. C., Gillah, A., Bates, E., Dec. 2007. Analysis of air quality within a street canyon using statistical and dispersion modelling techniques. Atmospheric Environment 41 (39), 9195–9205.
URL http://www.sciencedirect.com/science/article/pii/S1352231007006863
Wise, E. K., Comrie, A. C., Aug. 2005. Extending the Kolmogorov–Zurbenko Filter: Application to Ozone, Particulate Matter, and Meteorological Trends. Journal of the Air & Waste Management Association 55 (8), 1208–1216.
URL http://dx.doi.org/10.1080/10473289.2005.10464718
Zeileis, A., Kleiber, C., Krämer, W., Hornik, K., Oct. 2003. Testing and dating of structural changes in practice. Computational Statistics & Data Analysis 44 (1–2), 109–123.
URL http://www.sciencedirect.com/science/article/pii/S0167947303000306

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

Ziegler, A., König, I. R., Dec. 2013. Mining data with random forests: current options for real–world applications. WIREs Data Mining and Knowledge Discovery 4 (1), 55–63.
URL https://doi.org/10.1002/widm.1114

Appendix

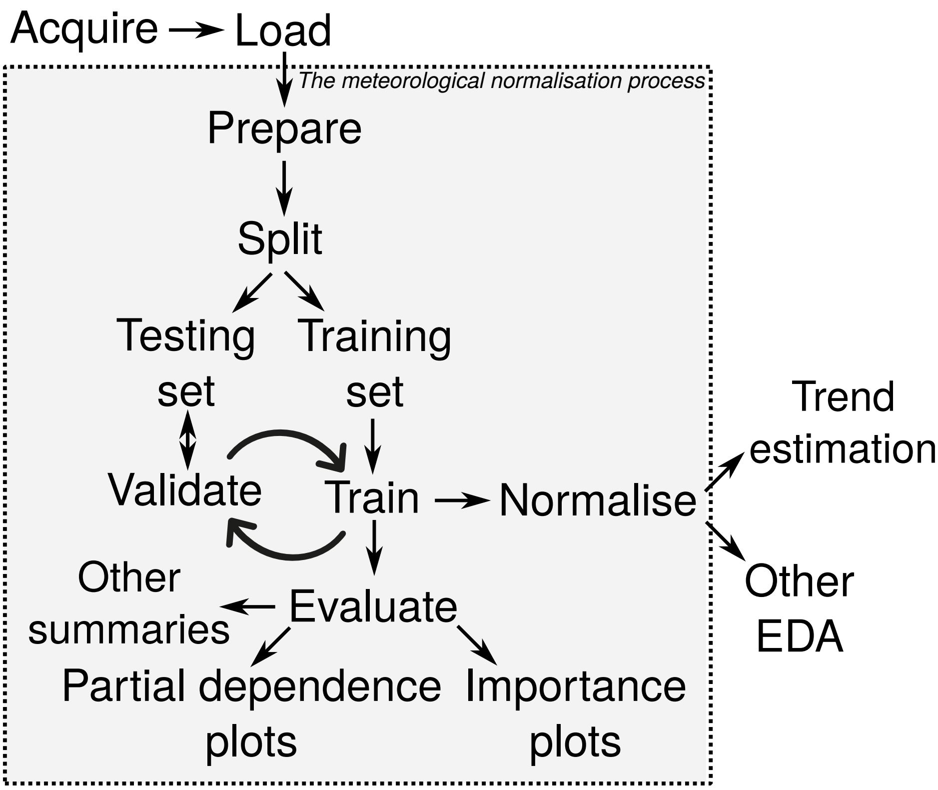

Figure A1: The framework for the meteorological normalisation technique. The training and validation phase is iterative to ensure the model does not overfit and adequate performance is achieved. After the technique has been completed, other analyses are conducted on the normalised time series.

Figure A2: Daily NO_2 and NO_x concentrations at London Marylebone Road between 1997 and 2016.

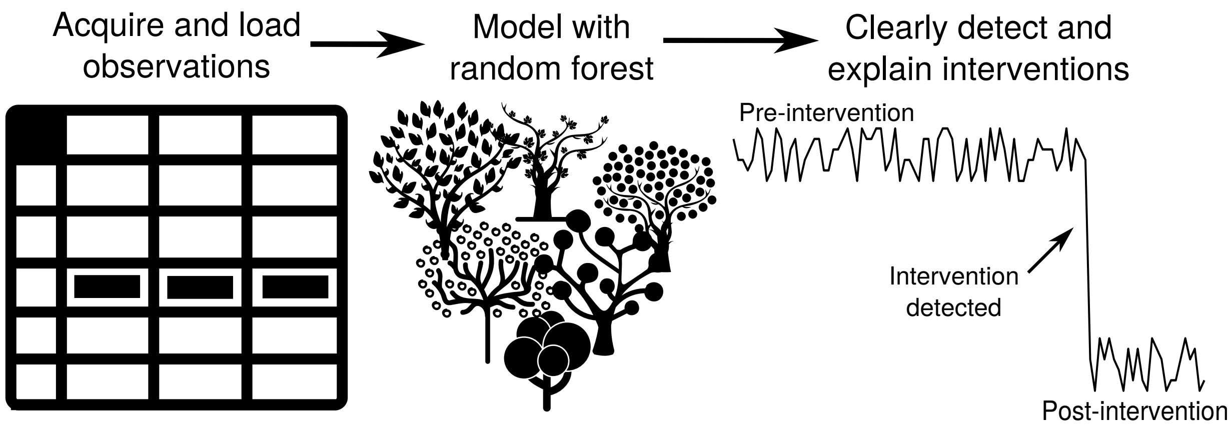

Figure A3: Graphical abstract. Icons designed by `freepik.com` from Flaticon.