

New Approaches to the Statistical Analysis of Air Quality Network Data

Insights from Application to National and Regional UK Networks

Polly Lang, BSc.



A thesis presented for the degree of
Doctor of Philosophy to the University of York

Wolfson Atmospheric Chemistry Laboratories
Department of Chemistry
University of York
September 2020

Abstract

Every hour, ambient concentration data for dozens of air pollutants is collected from hundreds of monitoring sites across the UK, adding to a repository consisting of more than 370 million observations going back 47 years. And yet, due to the difficulty of extracting meaningful information from this data, it is principally used for monitoring compliance with air pollution limits. This thesis aims to develop new statistical techniques and apply them to the air quality network data to derive additional insights concerning changes in air quality in the UK over the last twenty years.

The rolling change method is a new way of conducting robust long term trend analysis across multiple sites within monitoring networks that are subject to biases caused by site flux. It is used to analyse the long term trends in NO_x and NO_2 concentration, and in NO_2/NO_x ratio in London, Scotland, and the UK between 2000 and 2017. At each scale, the results are consistent, showing declines in NO_x and NO_2 concentration, and a peak in NO_2/NO_x around 2010, followed by a fall.

The 'meteorological normalisation' method using random forest is applied to remove the effect of meteorology from air pollutant concentrations at London sites to enable clearer visualisation of the trend due to changes in emissions. The method is also used to evaluate the efficacy of the London Low Emission Zone through the generation of counterfactual scenarios that are compared to the true normalised trend. The results suggest a mild improvement in air quality.

The influence of inter-annual meteorological variation on annual average concentrations of NO_x , NO_2 and O_3 is estimated for a large number of UK sites using the novel tools of heatmaps and cumulative sum plots. This influence is shown to be considerable: the range of the annual average concentration due to meteorological variation is $2.9 \mu\text{g m}^{-3}$ (8.2%) for NO_2 , $9.9 \mu\text{g m}^{-3}$ (12.6%) for NO_x and $3.3 \mu\text{g m}^{-3}$ (7.5%) for O_3 . The implications of these findings for the use of the annual average metric in compliance monitoring within the EU are considered.

Declaration

I declare that this thesis is a presentation of original work and I am the sole author. This work has not previously been presented for an award at this, or any other, University. All sources are acknowledged as References.

Chapter 2 is based on a peer-reviewed publication on which I was lead author. The full publication is available at the following reference:

Lang, P. E., Carslaw, D. C., & Moller, S. J. (2019). A trend analysis approach for air quality network data. *Atmospheric Environment: X*, 2, 100030.

Acknowledgements

I thank the Department of Chemistry at the University of York for funding this work.

I would also like to acknowledge Defra, King's College London and Ricardo Energy & Environment for providing most of the data used in this work, and Dr. Stuart Grange for allowing me to use the European air quality data from the smonitor database. Thank you to everyone who is involved with the enormous amount of effort that goes into maintaining the AURN, LAQN, SAQN, WAQN networks, and the European air quality monitoring sites and databases.

The Doctor of Philosophy programme presented in this thesis was not without its rocky moments. At times it hung by a thread, and it is with both a little surprise and a great deal of gratitude that I offer the following thanks.

I would like to thank my supervisors, Dr David Carslaw and Dr Sarah Moller for their unfailing support and guidance, and for not only sharing their ideas and insights with me, but also encouraging me to present my work at seminars and conferences. I am grateful too for their generosity in giving me the time and space to recover myself, and, later, the confidence to complete this PhD.

Thank you to all of the staff and students at the Wolfson Atmospheric Chemistry Laboratories, and in particular Dr. Stuart Grange, for being a friend, a teacher and an inspiration. Thank you Stuart, for being there for me when I really needed it.

Special thanks to Dr. Christoph Hüglin and his colleagues at EMPA, in Zurich for their hospitality and understanding.

And finally, to my family, who gave me life not once but twice: thank you.

Contents

1	Introduction	17
1.1	Background	18
1.1.1	Sources of air pollutants	19
1.1.2	Key chemistry of air pollutants	21
1.2	Measurement of ambient air quality in UK monitoring networks	25
1.2.1	NO _x measurement	26
1.2.2	PM measurement	27
1.2.3	O ₃ measurement	29
1.3	Air quality legislation and limits	29
1.4	Statistical analysis of ambient air quality data	30
1.4.1	Challenges associated with statistical analysis of ambient air quality data	31
1.4.2	Trend Analysis	33
1.4.3	Intervention Analysis	36
1.4.3.1	Air quality interventions	36
1.4.3.2	Analysing the impact of interventions	36
1.4.4	Compliance Monitoring	41
1.5	Structure and Content of this Thesis	43
1.5.1	Objectives	43
1.5.2	A new method for trend analysis in biased monitoring networks	44
1.5.3	Robust analysis of trends and policy interventions using random forest	45
1.5.4	Quantification of the effect of interannual meteorological variation on air pollutant concentration, and the implications for compliance metrics	45
2	A new trend analysis approach for air quality network data	54
2.1	Introduction	55

2.1.1	Background	55
2.2	Method	57
2.2.1	Identification of bias effects on the trend	57
2.2.2	Extraction of the underlying trend	58
2.2.3	Description of data	62
2.3	Results and Discussion	63
2.3.1	Testing the rolling change method through simulations	63
2.3.2	Long term trends in London ambient air quality	66
2.3.2.1	Identification of the bias effect on the trend	66
2.3.2.2	Extraction of the underlying trend	67
2.3.3	Long term trends in ambient air quality in Scotland and the UK	73
2.3.4	Potential applications	79
2.4	Conclusions	80
3	Development and Application of Random Forest Models to Air Pollutant Time Series	85
3.1	Introduction	86
3.1.1	Background	86
3.1.2	Intervention and Accountability Studies	86
3.1.3	Random Forest Modelling	86
3.1.4	Model Interpretation	88
3.1.4.1	Variable Importance	88
3.1.4.2	Partial Dependence Plots	89
3.1.5	Meteorological Normalisation	90
3.1.6	Chapter Summary	90
3.2	Methods	91
3.2.1	Data Preparation	91
3.2.2	Model Optimisation, Training and Evaluation	92
3.2.3	Trend Normalisation and Analysis	93
3.3	Results and Discussion	94
3.3.1	Method Validation	94
3.3.2	Method Analysis	96
3.3.2.1	Choice of meteorological data source	96
3.3.2.2	Normalised Trend Estimation using Partial Dependence Plots	99
3.3.3	Long Term Normalised Trends in London	101
3.3.4	Intervention analysis of the London Low Emission Zone	106

3.4	Conclusions & Next Steps	109
3.4.1	Conclusions	109
4	Quantifying the Effect of Inter-annual Meteorological Variation on Pollutant Concentrations	114
4.1	Introduction	115
4.1.1	Meteorology and Air Quality	116
4.2	Methods and Data	117
4.2.1	Description of Data	117
4.2.2	Modelling the Effect of Meteorology on Air Quality	119
4.2.3	Quantification of Variation	120
4.2.4	Visualisation of The Effect of Inter-annual Meteorological Variation on Air Quality Across the Monitoring Network	120
4.2.4.1	Heatmaps	121
4.2.4.2	Cumulative Sums	121
4.3	Results and Discussion	123
4.3.1	Quantifying the Effect of Inter-annual Meteorological Variation	125
4.3.1.1	NO ₂	125
4.3.1.2	NO _x	130
4.3.1.3	O ₃	133
4.3.2	The Effect of Inter-annual Meteorological Variation on Air Quality across the UK Monitoring Network	137
4.3.3	The Relationship between Ambient Concentration and Annual Meteorology	152
4.4	Conclusions and Future Directions	155
4.4.1	Implications for Compliance Monitoring and Air Quality Modelling	156
4.4.2	Future Directions	157
5.1	Future Directions	164
5	Conclusions & Future Directions	160
A	Appendix I: Algorithm for Chapter 2	167
B	Appendix II: Supplementary material for Chapter 2	169
C	Appendix III: aqtrends: An open source R package for air quality trend analysis	182

D Appendix IV: Publication	193
E Appendix V: Supplementary material for Chapter 4	211

List of Figures

1.1	Schematic representation of the OH catalysed oxidation of organic compounds, with ozone produced as a by-product. Image taken from the fourth report of the PORG (1997)	24
1.2	Extent of the London LEZ. Taken from Ellison et al. (2013)	37
1.3	The accountability chain (Henneman et al., 2017). Confounding factors that may affect each link of the chain are shown in red. This diagram is from the paper by Henneman et al. (2017)	37
2.1	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.	61
2.2	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 Appendix II (Table B.1).	63
2.3	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 1th 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.	65

2.4	Comparison of the average trend and rolling change trend ($n = 3$) with the true trend of data from 4 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 1th 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.	66
2.5	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 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 fit. The numbers at each data point correspond to the number of monitoring sites contributing to the data point.	68
2.6	Trend in UK NO_x emissions from road transport (urban driving) sectors between 2000-2016 (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 fit. . . .	70
2.7	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 NO_2 roadside increments at London roadside monitoring sites between 2005-2009 and 2010-2014. The error bars represent 95% confidence intervals.	72
2.8	The number of roadside monitoring sites measuring NO_x , NO_2 and PM_{10} concentration in Scotland by year between 2000 and 2017.	73
2.9	The spatial distribution of roadside monitoring sites measuring roadside NO_2 concentration in Scotland (a) at any point between between 2000 and 2017 (left), and (b) constantly throughout the duration of the period 2000-2017 (right).	74

2.10	The trends in roadside concentration of NO ₂ in Scotland between 2000 and 2017 calculated using three different methods: (i) the average concentration across all available data (top left), (ii) the average concentration using data only from sites measuring constantly throughout the analysis period (top right), and (iii) the rolling change method (bottom). The smoothed lines are loess (local regression) fits, with the 95% confidence interval represented by the shaded band. The numbers signify the number of monitoring sites contributing to each annual data point.	75
2.11	Annual rolling change trends in NO _x , NO ₂ , NO ₂ /NO _x and PM ₁₀ concentrations in Scotland 2003-2018. The numbers indicate the number of monitoring sites contributing to each (annual) data point.	76
2.12	Annual rolling change trends in NO _x and NO ₂ concentration and NO ₂ /NO _x ratio at roadside in the UK 2000-2017. The numbers indicate the number of monitoring sites contributing to each (annual) data point.	78
2.13	Annual rolling change trends in NO _x and NO ₂ concentration and NO ₂ /NO _x ratio at roadside in Europe 2000-2017. The numbers indicate the number of monitoring sites contributing to each (annual) data point.	78
3.1	Monthly average normalised trend in black carbon concentration at Marylebone Road, London 2008-2018. The vertical dashed lines indicate the break points (as calculated using the regression model break point analysis described by Bai (1994)).	95
3.2	Comparison of the normalised trend in PM ₁₀ concentration at (a) London Marylebone Road (MY1), (b) Cromwell Road (KC2), and (c) Camden Swiss Cottage (CD1) 2000-2017, calculated using models trained on NOAA surface meteorological measurement data 'noaa', hourly ethane concentration measured at MY1 ('ethane'), and on no meteorological data at all ('none').	98
3.3	Comparison of the normalised trend and the partial dependence plot for the Unix date (long term trend trend) variable from the random forest model of PM ₁₀ concentration at Marylebone Road, London 2000-2017.	100

3.4	Long term trends in (a) PM_{10} , (b) NO_x , and (c) NO_2 concentration at, from left to right, Marylebone Road (MY1), Camden Kerbside (CA1), and Cromwell Road (KC2) in London between 2000 and 2017. The purple line is the partial dependence plot of the date variable calculated from the random forest model trained using monitoring data from the specified site. The vertical dashed lines indicate the break points.	102
3.5	Comparison of the normalised trend in PM_{10} concentration (purple line) at (a) Marylebone Road, (b) Camden Kerbside, and (c) Cromwell Road 2000-2017 with the proportion of the UK vehicle fleet composed of post-Euro 5 vehicles over the same period (yellow bars). The vertical dashed lines indicate the break points.	105
3.6	Comparison of the normalised trends in PM_{10} concentration at Marylebone Road, London 2008-2017 with the intervention indicator variable set to the values for Phases 0-3. These trends represent different counterfactual scenarios (i.e. the observed trend had each phase of the LEZ been in place for the entire time period).	107
3.7	Summary of the results of the London LEZ intervention analysis (the effect of the London LEZ on ambient PM_{10} concentrations) carried out using data from nine London monitoring sites. The segments indicate the average difference in PM_{10} concentration between normalised trends for two different phases (values of the intervention indicator variable), with red indicating a higher PM_{10} concentration for the higher phase, blue indicating a lower concentration, and green indicating no significant change (paired t-test).	108
4.1	Distribution of monitoring sites measuring (4.1a) NO_2 , (4.1b) NO_x , and (4.1c) O_3 concentrations with at least 80% data capture between 2008 and 2017 in the UK.	119
4.2	Meteorologically normalised time series for (4.2a) NO_2 , (4.2b) NO_x , and (4.2c) O_3 concentration 2008-2017 at Marylebone Road, for each meteorological year. The plot on the left shows the trends for each meteorological year 2008-2017, and the plot on the right shows the range resulting from meteorological variation (the range of predicted concentrations over all meteorological years) as an envelope around the average normalised trend.	125

4.3	Range in the annual average NO ₂ concentration at all UK sites in each year due to meteorology. The data points show the mean NO ₂ concentration in that year at an individual monitoring site, while the error bars represent the range in the mean value resulting from meteorological variation. The range was calculated as the range in the predicted concentration for that year over normalised data from all meteorological years. The plot on the right is a zoomed in view of the plot of the left, showing the ‘marginal’ sites (those whose compliance with the EU limit values falls within the range due to meteorological variation) in more detail.	128
4.4	Spatial distribution of the sites which are ‘marginal’ for NO ₂ (i.e. their compliance with EU limit values falls within the range (min to max) due to inter-annual meteorological variation in at least one year between 2008 and 2017).	129
4.5	Annual average NO ₂ concentrations, and range (min to max) due to meteorological variation (represented by the error bars), averaged across all UK sites over the period 2008 to 2017.	130
4.6	Range in the annual average NO _x concentration at all UK sites in each year due to meteorology. The data points show the mean NO _x concentration in that year at an individual monitoring site, while the error bars represent the range (min to max) in the mean value resulting from meteorological variation. The range was calculated as the range in the predicted concentration for that year over normalised data from all meteorological years.	132
4.7	Annual average NO _x concentrations, and ranges due to meteorological variation (represented by the error bars), averaged across all UK sites over the period 2008 to 2017.	133
4.8	Range in the annual average O ₃ concentration at all UK sites in each year due to meteorology. The data points show the mean O ₃ concentration in that year at an individual monitoring site, while the error bars represent the range of the mean value resulting from meteorological variation. The range was calculated as the range in the predicted concentration for that year over normalised data from all meteorological years.	136
4.9	Annual average O ₃ concentrations, and ranges due to meteorological variation (represented by the error bars), averaged across all UK sites over the period 2008 to 2017.	137

4.10	Heatmaps showing the difference between the mean concentration in each meteorological year and the average concentration over all meteorological years for (4.10a) NO ₂ , (4.10b) NO _x , and (4.10c) O ₃ , for each UK monitoring site, as predicted by the random forest model.	140
4.11	Distribution of monitoring sites with unusual NO _x and NO ₂ heatmap results (unusual variation in NO _x and NO ₂ concentration as a result of inter-annual meteorological variation).	142
4.12	Cumulative sum (CUSUM) plots of the difference between the mean concentration in a given meteorological year and the average concentration over all meteorological years for NO ₂ , NO _x and O ₃ (columns 1-3 of each grid). The sites are ordered by (4.12a) mean pollutant concentration, (4.12b) NO ₂ /NO _x ratio, (4.12c) latitude (north to south from left to right), and (4.12d) longitude (east to west from left to right). Note: In the case of the O ₃ CUSUM plots ordered by NO ₂ /NO _x ratio, only sites also measuring NO ₂ and NO _x are used, and they are ordered by the NO ₂ /NO _x ratio derived from their measurements.	148
4.13	Summary of the results of the CUSUM analysis for NO ₂ , NO _x and O ₃ - the Theil-Sen slope of the CUSUM plot (ordered by mean pollutant concentration), which represents the average by-site deviation of the concentration due to meteorology from the average concentration 2008-2017. The meteorological years are ordered by their effect on the ambient NO ₂ concentration, from the year which caused the largest decrease in NO ₂ concentration on the left, to the year whose meteorology resulted in the greatest elevations of NO ₂ concentration on the right.	151
4.14	A comparison of the average UK meteorological conditions (air temperature, air pressure, relative humidity, dew point, wind speed, wind direction and visibility) as a function of the Julian day in each year 2008–2017 (purple smooth lines) and the average UK meteorological conditions over the entire ten-year period (yellow smooth lines). All hourly data used as predictors during random forest model training were averaged to daily data, and the line represents a loess smooth through the daily averages. The shaded band represents the 95% confidence interval around the loess smooth.	153

B.1 Comparison of the rolling change trend of simulated data calculated using window widths $n = 2, 3,$ and 5 . The trend was calculated from 100 simulated time series, which were randomly sampled 100 times from the ‘combined’ scenario. The lines correspond to the trends with NCC equal to the 50th, 10th and 1th 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. 172

B.2 Difference (left), cumulative sum of differences (centre) and cumulative sum of differences weighted by the number of sites opening and closing (right) in between the average concentration of opening sites and closing sites in each year for (a) NO_x , (b) NO_2 , and (c) NO_2/NO_x at London roadside sites 2000-2017. The lines represent a loess smooth fit to the data, and the shaded bands represent the 95% confidence interval around the smooth. 174

B.3 Rolling trends (left) and average trends (right) in NO_x concentration for $n =$ (a) 3, (b) 5, (c) 10, and (d) 12 at London roadside sites 2000-2017. The average trend (right) was calculated using data from the same sites as the rolling trends, which are filtered by site duration based on the value of n . The black lines in the right hand plots represent a loess smooth fit to the data, and the shaded bands represent the 95% confidence interval around the smooth. 176

B.4 Rolling trends (left) and average trends (right) in NO_2 concentration for $n =$ (a) 3, (b) 5, (c) 10, and (d) 12 at London roadside sites 2000-2017. The average trend (right) was calculated using data from the same sites as the rolling trends, which are filtered by site duration based on the value of n . The black lines in the right hand plots represent a loess smooth fit to the data, and the shaded bands represent the 95% confidence interval around the smooth. 177

B.5 Rolling trends (left) and average trends (right) in NO_2/NO_x concentration for $n =$ (a) 3, (b) 5, (c) 10, and (d) 12 at London roadside sites 2000-2017. The average trend (right) was calculated using data from the same sites as the rolling trends, which are filtered by site duration based on the value of n . The black lines in the right hand plots represent a loess smooth fit to the data, and the shaded bands represent the 95% confidence interval around the smooth. 178

B.6 Trends in NO_x at London urban background sites 2000-2017 using data from all available monitoring sites (left) and data from long term monitoring sites only (right). 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 average. 179

B.7 Difference (left), cumulative sum of differences (centre) and cumulative sum of differences weighted by the number of sites opening and closing (right) in between the average concentration of opening sites and closing sites in each year for NO_x at London urban background sites 2000-2017. The lines represent a loess smooth fit to the data, and the shaded bands represent the 95% confidence interval around the smooth. 180

List of Tables

1.1	European Directive air quality limit values for selected air pollutants for the United Kingdom (Defra, 2019).	30
1.2	European Directive limit values of the concentration of selected air pollutants (European Commission, 2019)	42
2.1	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 , NO ₂ and NO ₂ /NO _x concentration at roadside in London 2000-2017.	69
4.1	Random forest model performance metrics (averaged over all monitoring sites). The root mean squared errors (RMSE) are given in µg m ⁻³	123
4.2	Number of sites with unusual heatmap results for NO _x and NO ₂ , by region and site type.	142
B.1	Metadata for London roadside monitoring sites used in the analysis. The mean and standard deviation of the hourly NO _x and NO ₂ concentration data measured by each site during the period of 2000 to 2017 is also shown. Monitoring sites were selected as those within the bounding box of coordinates 51.25°N, 51.71°N, -0.54°E, 0.28°E. The data was sourced from the AURN, LAQN and AQE networks.	170
E.1	Metadata for the monitoring sites used in the analysis. The data was sourced from the AURN, LAQN and AQE networks.	212

1. Introduction

1.1 Background

Air pollution is increasingly being recognised as one of the most pressing problems facing the world today. Its impact is felt in a wide range of contexts including economies, agriculture, ecosystems, infrastructure and health, contributing to an estimated 40,000 premature deaths a year in the UK alone and costing the economy more than £20 billion per year (Royal College of Physicians, 2016).

Exposure to air pollution contributes to a myriad of health effects, including heart disease, strokes and pulmonary disease. The exact mortality due to air pollution is difficult to quantify, because of the complicated relationship with other risk factors. However, air pollution is considered to be the fourth greatest cause of mortality worldwide, killing more than six times as many people as malaria and four times as many as HIV/AIDS every year (The World Bank and Institute for Health Metrics and Evaluation, 2016).

Thousands of pollutants contaminate the atmosphere, but of particular concern are the concentrations of nitrogen oxides (NO_x), which are composed of NO and NO_2 , particulate matter (PM), and tropospheric ozone (O_3).

Nitrogen oxides (NO_x) are defined as the sum of nitrogen monoxide (NO) and nitrogen dioxide (NO_2) (Schultz et al., 2015). NO_x is emitted by both natural sources, such as bacterial production, volcanic activity and lightening, and anthropogenic sources, primarily fossil fuel and biofuel combustion (e.g. vehicle exhausts, power stations, etc.) (Schultz et al., 2015). Natural sources and secondary production generate a background NO_2 concentration, while primary anthropogenic emissions cause spikes in NO_x and NO_2 concentration above the background near sources. While NO typically dominates overall NO_x concentrations in anthropogenic emissions, in the atmosphere NO reacts quickly with O_3 to form NO_2 , therefore in well-mixed background air (distant from an anthropogenic source), the ratio of NO to NO_2 concentration (NO/NO_2) is lower (Carslaw et al., 2016).

Particulate matter is a broad categorisation describing all airborne particles. Composition and size vary widely, with particles made up of a complex mixture of organic and inorganic species and ranging in diameter from less than 0.1 μg to 100 μg . Two main components of PM are currently measured at monitoring stations for compliance monitoring. PM_{10} refers to particles with a diameter of < 10 μg , which are small enough to be inhaled through the nose to enter the body. $\text{PM}_{2.5}$ (the 'fine fraction') includes particles with a diameter of < 2.5 μg (Anderson et al., 2012; Adams et al., 2015). These fine particles are small enough

to penetrate to the lungs and into the bloodstream. Particles larger than 10 μg are not regulated because they do not enter the lungs. On the other hand, fine particles that can enter the lungs and bloodstream can cause adverse health effects (Adams et al., 2015). Background PM comes from natural sources, such as volcanic activity and dust storms, dispersion, mixing and secondary production (Anderson et al., 2012). Anthropogenic sources include fossil fuel combustion, biomass combustion, agriculture, industrial processes and uplifting of dust into the atmosphere (Adams et al., 2015; Anderson et al., 2012).

Ozone is a secondary pollutant, and is mostly produced through photochemical reactions in the atmosphere (Wang et al., 2019). O_3 exists in a photochemical equilibrium with NO and NO_2 , and volatile organic compounds (VOCs) or hydroxyl radicals. The position of the equilibrium depends on the relative concentrations of the precursors NO_x and VOCs. Higher and lower ratios can favour ozone depletion, therefore the relationship between O_3 concentration and anthropogenic emissions is complex (Wang et al., 2019; Monks, 2005). O_3 concentration is highly dependent on dispersion conditions and long-range transport. High concentrations of O_3 are damaging to human health, agriculture and environments (Schultz et al., 2015; Wang et al., 2019).

An understanding of the sources and the atmospheric chemistry of these air pollutants is vital to interpreting the results of air quality data analysis. The following sections describe the emission sources of the major air pollutants and the reactions they undergo once emitted into the atmosphere.

1.1.1 Sources of air pollutants

Nitrogen oxides (NO_x) are defined as the sum of nitrogen monoxide (NO) and nitrogen dioxide (NO_2) (Schultz et al., 2015). Most NO_x is emitted from the combustion of fuels, with other sources making up just 3% of total emissions in 2018 (NAEI, 2020). Emissions of NO_x from combustion come from a variety of fuels and sectors. In 2018, 31% of NO_x emissions came from road transport, 21% from other forms of transport, 20% from energy production (e.g. power stations) and 8% from other industrial activity (NAEI, 2020).

NO_x is mainly emitted as NO , which is rapidly converted to NO_2 through the reaction with O_3 (see Section 1.1.2), however some NO_2 is emitted from primary sources, particularly road vehicles (AQEG, 2004).

During fuel combustion, NO_x is produced either by the high temperature oxidation of nitrogen from the air, or from nitrogen chemically present in the fuels themselves. NO_x formation is favoured by high temperature and oxidation rich

conditions (AQEG, 2004).

The main mechanism by which NO is formed is known as the Zel'dovich mechanism (the thermal-NO mechanism) which is shown in Equation 1.1 (AQEG, 2004).



Only an estimated 5% of NO_x emitted from combustion is released as NO₂, but the ratio of NO₂/NO emitted from diesel engines is higher than from gasoline. NO₂ production is favoured in poorly mixed systems and by low temperatures. NO is converted to NO₂ via the reaction shown in Equation 1.2 (AQEG, 2004).



Some NO₂ is also converted back to NO according to Equation 1.3 (AQEG, 2004).



The total concentration of particulate matter (PM) is made up of a primary component, which is directly emitted from primary sources, as well as a large secondary component, which is formed from the chemical reactions of gases and other aerosols in the atmosphere.

PM₁₀ is particulate matter with a diameter of less than 10 μg (Anderson et al., 2012; Adams et al., 2015). In 2018, in the UK, 40% of PM₁₀ emissions came from industrial, commercial or residential combustion, 32.5% from production processes, 12% from agriculture, 11% from road transport, 3% from other transport and 1.5% from public electricity and heat production (NAEI, 2020).

PM_{2.5} is particulate matter with a diameter of less than 2.5 μg (Anderson et al., 2012; Adams et al., 2015). In 2018, in the UK, 64% of PM_{2.5} emissions came from industrial, commercial or residential combustion, 12% from production processes, 11% from road transport, 6% from agriculture, 5% from other transport, and 2% from public electricity and heat production (NAEI, 2020).

Vehicles produce exhaust PM, from fuel combustion, and non-exhaust PM, which is generated by mechanical abrasion and corrosion from tyre, brake, clutch and road surface wear and from vehicle component corrosion. Road transport

also causes suspension or re-suspension of road dust (AQEG, 2005, 2019).

Natural sources also contribute to primary PM emissions. These include wind-suspended dust, sea spray, forest fires, volcanic activity and biogenic emission (AQEG, 2012; Anderson et al., 2012).

Unlike the other air pollutants discussed here, ozone (O_3) is a secondary pollutant formed by photochemical oxidation of volatile organic compounds (VOCs) in the presence of nitrogen oxides (NO_x) (Wang et al., 2019). The details of this reaction are discussed further in Section 1.1.2.

Although it is not emitted directly, the production of ozone depends on the concentration of its precursors, which are primary pollutants. The sources of NO_x have already been described. In 2018 in the UK, the largest contribution to anthropogenic NMVOC (non-methane VOC) emissions came from industrial processes and product use, which accounted for 53% of total emissions. Extraction and distribution of fossil fuels accounted for 19%, and agriculture for 12%. 7% of emissions came from transport and other mobile sources (exhaust emissions and evaporative losses of fuel vapour), and 9% from stationary combustion sources (NAEI, 2020).

VOCs are also emitted from vegetation and the soil. Emissions from natural sources are dependent on temperature, light levels and plant species and therefore have strong seasonal variation. Most biogenic VOC sources are not normally co-located with major urban NO_x sources, whereas anthropogenic sources of VOCs are more often co-located with NO_x sources in urban areas. Therefore, in urban areas where high concentrations of the precursors result in high ozone production, anthropogenic sources of VOCs may be relatively more important. However, over recent years, as anthropogenic VOC emissions have been reduced, the relative importance of biogenic VOCs has increased, a trend that is expected to continue in the future (AQEG, 2020).

1.1.2 Key chemistry of air pollutants

Once emitted into the atmosphere, air pollutants are subject to dispersion and reaction with other atmospheric species. Since both of these processes are highly dependent on meteorology, the meteorological conditions exert a strong influence on ambient air quality (AQEG, 2004).

The ambient concentration of air pollutants is dependent on the combination of (i) emissions, (ii) chemistry, and (iii) meteorology (AQEG, 2004). It is therefore crucial during interpretation of results to understand all three factors and their

effects on the ambient concentrations of the air pollutant under investigation.


In particular, the interconversion between NO and NO₂ is crucial to an understanding of both NO_x and ozone concentrations.

While most NO_x is emitted as NO, once in the atmosphere NO reacts rapidly with ozone to form NO₂ via the reaction shown in Equation 1.4.



In locations with high concentrations of NO, such as polluted roadside locations, the timescale for conversion of NO to NO₂ via this mechanism can be as little as 2 seconds (AQEG, 2004).


During the day, NO is regenerated from the photolysis of NO₂ via the reaction shown in Equation 1.5 (AQEG, 2004).



The rate of NO₂ photolysis varies with season, time of day and latitude, but is slower than the conversion of NO to NO₂. The average lifetime of NO₂ in the summer is about 3 minutes, and two or three times longer in winter (AQEG, 2004). Even so, the daytime interconversion of NO and NO₂ is so efficient that the species are usually referred to in combination as NO_x.


As previously mentioned, the majority of NO_x is emitted as NO, but once in the atmosphere the NO reacts rapidly with ozone to form NO₂. The two species subsequently interconvert. The result of this is that the NO₂/NO_x ratio is at a minimum close to the emission source, and increases with distance until a stable background NO₂/NO_x ratio is reached in well-mixed air. The NO₂/NO_x ratio, therefore, can be used as a proxy for the degree of contribution to the pollutant concentration at a particular location, and can be used to distinguish between 'traffic' locations (close to road traffic sources) and 'background' locations. This utility is applied in Chapter 4.

The main loss mechanism for NO_x is conversion of NO_x to nitrate, followed by wet or dry deposition or further reaction. The reaction is shown in Equation 1.6 (AQEG, 2004).





During the night, nitrate is also generated from NO_x via the route shown in

Equation 1.7




In addition to deposition, nitrate is also removed via formation of nitrate aerosol, or by reaction with gaseous ammonia to form ammonium nitrate aerosol.


NO_x is also removed temporarily by formation of the reservoirs nitrous acid (HONO) and PAN. These reactions are shown in Equations 1.8 and 1.9 (AQEG, 2004).



During the day, NO is regenerated from photolysis of HONO. PAN exists in an equilibrium with NO_2 , and undergoes thermal decomposition under conditions characterised by high temperatures and high levels of organic radicals (AQEG, 2004).

Ozone is formed photochemically by the sunlight-initiated oxidation of VOCs in the presence of NO_x . As previously shown in Equation 1.5, photolysis of NO_2 produces NO and O_3 . Since, as shown in Equation 1.4, NO and O_3 react to re-form NO_2 , these reactions constitute a cycle with no net ozone produced. However, in the presence of volatile organic compounds, the oxidation of the organic species catalysed by OH radical consumes NO and produces NO_2 . The NO_2 photolyses to re-form the NO, which is used in further VOC oxidation. Ozone is produced as a by-product of the cycle (Monks, 2005). The reactions involved in ozone formation are shown in Equation 1.10 (PORG, 1997).





The series of reactions is perhaps more clearly visualised diagrammatically as a cycle, as shown in Figure 1.1 (PORG, 1997).




Figure 1.1: Schematic representation of the OH catalysed oxidation of organic compounds, with ozone produced as a by-product. Image taken from the fourth report of the PORG (1997).


The relationship between the production of ozone and the ambient concentrations of NO_x and VOCs is not linear, but rather depends upon the relative concentrations of each precursor. In VOC-limited environments, where the relative concentrations $[\text{NO}_x]/[\text{VOC}]$ is high, an increase in VOC emissions would lead to increased ozone production, but an increase in NO_x emissions would decrease the rate of ozone production. The reverse situation presents itself in NO_x -limited environments, where $[\text{NO}_x]/[\text{VOC}]$ is low. Here, an increase in NO_x emissions will increase ozone production, while an increase in VOC emissions will cause decreased ozone production (Monks, 2005; PORG, 1997).

The reason behind this lies in the cyclical nature of the process. The rate of ozone production is proportional to the chain length, which is the number of VOC-oxidation cycles (as shown in Figure 1.1) which occur before radical removal. Changes in NO_x and VOC concentrations affect the rate of the chain terminating

reactions relative to the competing reactions (Monks, 2005; PORG, 1997).

Secondary PM is formed by gas-to-particle conversion of the lower volatility or soluble products or intermediates that arise from oxidation of higher volatility emitted gases, or by the reaction of gases with aerosols or aqueous droplets.

One example of this is the formation of sulfate aerosol from SO_2 . SO_2 can be oxidised to sulfate in the gaseous phase or the aerosol phase. Another important primary pollutant, NO_x , is also oxidised in the atmosphere to form nitrate aerosol. Nitrate is formed through the oxidation of NO_2 , as shown in Equation 1.6 (AQEG, 2005).



Sulfate and nitrate are often present in PM as ammonium sulfate and ammonium nitrate, which can form from ammonia emitted mainly from agricultural sources (AQEG, 2005; Adams et al., 2015). Oxidation of volatile organic compounds (VOCs) produces secondary organic aerosol (SOA) via a hugely complex variety of species and mechanisms (AQEG, 2005).

Because a considerable component of the total atmospheric PM comes from secondary sources, regional chemistry and transport exerts a powerful influence on the ambient concentration of PM. For example, while some PM in UK urban areas comes from primary particle emissions, it is frequently the case that the bulk comes from regional chemistry and European transport. For this reason, caution must be exerted during analysis not to misattribute changes in background PM concentrations to changes in the emission source. This becomes particularly relevant in Chapter 3, which deals explicitly with the identification of the background component of ambient PM concentration.


1.2 Measurement of ambient air quality in UK monitoring networks

Many techniques and instruments have been used to measure the ambient concentrations of air pollutants. Here, the methods used within the Automatic Urban and Rural Network (AURN) for continuous air quality monitoring are described. The majority of the data used in this thesis to analyse UK ambient air quality was sourced from AURN sites. However, it is important to note that the data within this thesis was sourced from multiple independent monitoring

networks and, in some cases, was measured using different instruments. This has implications for the interpretation of the results, as will be discussed.

1.2.1 NO_x measurement

NO_x and NO₂ are measured at AURN sites using chemiluminescent analysis. Within the sampler, NO reacts with O₃ to produce an excited state of NO₂ (see Equation 1.11). When the excited molecule of NO₂ returns to the ground state, radiation is emitted. The intensity of the radiation (fluorescence) is proportional to the concentration of NO present (AQEG, 2004).



The analyser does not directly measure NO₂. Instead, the NO₂ in the ambient air is reduced to NO in the NO₂-to-NO converter (usually a heated molybdenum converter in the UK) (Schultz et al., 2015). The analyser then independently measures the concentrations of NO and NO_x, and calculates the concentration of NO₂ that was originally present by subtracting NO from total NO_x (Equation 1.12). Some instruments achieve this by measuring NO_x and NO simultaneously in dual chambers (AQEG, 2004).

$$[\text{NO}_2] = [\text{NO}_x] - [\text{NO}] \tag{1.12}$$

Typically, chemiluminescent analysers have a lower detection limit of < 2 µg m⁻³. This is much lower than the alternative NO₂ measurement methods, such as diffusive samplers, electrochemical sensors, thick film sensors and differential optical absorption spectroscopy (DOAS), which have detection limits on the order of several µg m⁻³ (AQEG, 2004).

One problem with chemiluminescent analysis is interference from other reactive nitrogen species (Schultz et al., 2015). In the molybdenum NO₂-to-NO converter, some oxidised nitrogen species, such as HNO₃, HONO, PAN and alkyl nitrates, are converted to NO, therefore contributing to the measured NO₂ concentration. Only HONO and PAN are considered to interfere significantly under ambient conditions present in the UK. Under typical urban background conditions, the interference from HONO and PAN has been estimated to account for ≤ 2% and ≤ 5% of the measured NO₂ concentration respectively (AQEG, 2004).

In the UK, the Air Quality Directive specifies an upper limit for the uncertainty with which air quality measurements are made (European Union, 2008). The

upper limit uncertainty for continuous measurements of NO_x and NO_2 is 15%. This uncertainty may proceed from many factors related to the instrument itself, instrument calibration and instrument operation. It is important, when interpreting the analysis of data collected from these instruments, to recognise that the measurement are subject to this degree of uncertainty and avoid over-interpretation (AQEG, 2004).

1.2.2 PM measurement

In the UK, PM is typically measured as a mass concentration. There are a wide variety of methods in use for this.

Filter-based gravimetric samplers are the reference method of measurement defined in the EU First Air Quality Daughter Directive. A sampling inlet is connected to a filter substrate and a flow controller. PM_{10} is collected on the filter during the sampling period, then the mass determined gravimetrically. However, the time resolution of the measurements is limited to 24 hours, which limits its utility in detailed statistical analysis. It also has high operating costs, and cannot be used for continuous monitoring due to the high degree of human operation required (for example, daily filter exchange and weighing of filters). Human error leads to more potential pathways for error, and there is always a delay between sample collection and reporting, while the samples are transferred, stored and weighed. Consequently, the AURN monitoring sites use methods with demonstrated equivalence to gravimetric sampling to provide continuous real-time hourly monitoring data (AQEG, 2005).

TEOM (Tapered Element Oscillating Microbalance) analysers are widely used to measure continuous concentrations of PM in the UK. Air enters a tapered glass tube and is deposited onto a small filter at one end. The frequency of oscillation of such tubes is directly proportional to the mass of the tube, so the change in mass of the tube is proportional to the additional mass of the deposited PM (AQEG, 2005).

TEOM analysers not only impart greater precision than the reference method, but also provide real-time data with hourly resolution, which has greater utility in statistical analysis. Differences in moisture equilibration between TEOM and gravimetric samplers means that an adjustment factor must be applied to the measured data in order for the measurements to be equivalent (AQEG, 2005).

The FDMS (Filter Dynamic Measurement System) instrument is a development of the TEOM analyser which can account for semi-volatile PM. Data is collected hourly and daily. Water is removed from the air sample before it is passed into

the sampling unit, and the operating temperature is lower. After the air is passed through the filter to collect non-volatile PM, the air is diverted into a cooled chamber which is maintained at 4 °C and passed over a filter. The air is then returned to the first sensor unit to provide a baseline measurement (AQEG, 2012).

While TEOM and FDMS samplers are the most common method for continuous PM measurement in the UK, β -attenuation analysers are more widely used in Europe. This method measures the reduction in intensity of β particles passing through a filter, which is proportional to the mass of PM collected on the filter. Data is measured at hourly resolution, but the method has lower precision than the TEOM analyser (AQEG, 2005).

Less widely used methods for PM measurement include optical analysers, the black smoke method, and personal samplers (AQEG, 2005).

Measurement of particulate matter is particularly challenging and the choice of measurement method often has a considerable effect on the results. This is because different instruments may treat the air stream in different ways. For example, TEOM analysers preheat the air stream, while filter-based gravimetric samplers do not, which results in the greater loss of semi-volatiles from the former. Empirically determined adjustment factors are applied to make the data measured using each method equivalent, but for this reason, data measured using different methods may not be comparable (AQEG, 2005). There is also variation between instruments of the same type. The filter material can influence the measured data, because the filter mass changes in response to humidity and the accumulation of water or other material. The mass change can vary by the type of filter material, and even by difference batches of the same material, so this uncertainty is very difficult to control or account for (AQEG, 2012). Other reactive gases may also absorb onto the filter or the PM, contributing to further measurement uncertainty (AQEG, 2012).

All of these methods can be adapted to measure $PM_{2.5}$ as well as PM_{10} , for example, by changing the size of the fractioning inlet. However, there is a smaller amount of $PM_{2.5}$ in the ambient air than PM_{10} , which makes accurate $PM_{2.5}$ measurement more challenging than measurement of the coarser fractions (AQEG, 2005). In addition, the $PM_{2.5}$ particles generally contain a larger proportion of semi-volatile and hygroscopic material, which are more difficult to capture (AQEG, 2012).

Accurate measurement of PM has become more challenging as PM concentrations have decreased. Well known problems such as changes in mass of the filter due to water collection or loss of fibres become more significant when attempting

to measure smaller masses (AQEG, 2012).

1.2.3 O₃ measurement

UV spectroscopy is used to measure ambient ozone concentration at UK monitoring sites. The air sample passes through a cell of length l and the intensity of UV light at 254 nm (I_1) is measured by a detector. A zero reference intensity (i.e. the intensity of the air with no ozone present), I_0 , is calculated using air that has been passed through an ozone-removing scrubber. The concentration is then calculated using the Beer-Lambert law shown in Equation 1.13, where ε is the absorption coefficient at 254 nm (AQEG, 2009).

$$I_0 = I_1 e^{-\varepsilon cl} \quad (1.13)$$

The presence of other species that are optically active at 254 nm could potentially interfere with the ozone measurement, but in practise this is minimal, since the same species are present in both the reference and the non-reference air samples, and so any potential interference is cancelled out (AQEG, 2009).

The ozone analysers are calibrated on-site using ozone photometers, with an uncertainty of 3.5%. The ozone photometers are themselves calibrated against the UK ozone standard with an uncertainty of 3.0%, which is calibrated against other national measurements' standards. Overall, the maximum uncertainty of the ozone data is 12% (AQEG, 2009). While this is below the directive of 15%, it still merits consideration during data analysis and interpretation.

1.3 Air quality legislation and limits

Many governing bodies have attempted to improve air quality by reducing anthropogenic emissions of air pollutants. The European Union imposes limits on the concentrations of various air pollutants based on the World Health Organisation's ambient air quality guidelines. A selection of the EU limits are shown in Table 1.1. However, the limits are regularly breached in many of the member states, including the UK. In fact, in May 2018 the EU referred France, Germany and the UK to the Court of Justice of the EU for their repeated exceedances of the NO₂ concentration limit values and their failure to take appropriate measures to reduce the frequency and duration of these exceedances (Airqualitynews.com, 2018).

Table 1.1: European Directive air quality limit values for selected air pollutants for the United Kingdom (Defra, 2019).

Pollutant	Limit	Concentration measured as
NO ₂	200 $\mu\text{g m}^{-3}$ not to be exceeded more than 18 times a year	1 hour mean
NO ₂	40 $\mu\text{g m}^{-3}$	annual mean
PM ₁₀	50 $\mu\text{g m}^{-3}$ not to be exceeded more than 35 times a year	24 hour mean
PM ₁₀	40 $\mu\text{g m}^{-3}$	annual mean
PM _{2.5}	25 $\mu\text{g m}^{-3}$	annual mean

In order to control the air pollution, it is important to study how the concentrations of air pollutants have changed in the past, and consider the factors that drove the change. Only by evaluating changes and trends in air pollutant concentrations and their drivers can the effectiveness of interventions aimed at improving air quality be quantified. In this context, it is clear that accurate analysis of air pollution is vital, not merely to evaluate compliance with legislation, but to preserve economies, ecosystems and human health.

1.4 Statistical analysis of ambient air quality data

Historically, atmospheric modelling based on detailed emissions inventories has been the primary technique for predicting and analysing air pollution. This has many limitations however, namely, the reliance on the use of emissions factors and flux estimates, which may be inaccurate and have high uncertainty. Modelling also requires detailed understanding of all of the processes and interactions of the system which, in the case of the atmosphere, are often complex or poorly understood. The complex nature of atmospheric processes and local conditions means 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 actual pollutant concentrations, correlations and trends at measurement locations. Such information is invaluable for determining the actual effects of social and infrastructure changes and policy interventions on air quality.

The direct quantification of the actual effects of interventions using statistical analysis of ambient air quality data is very attractive to policy makers, though many factors conspire to make it a complex undertaking, as shall be discussed in Section 1.4.1. Air quality data analysis aims to identify sources of air pollutants,

which enables targeting of the most relevant sources in policy interventions, as well as identifying trends in pollutant concentrations, and interpreting them in terms of events e.g. interventions.

The UK air quality monitoring network has been in place since 1973, and has collected more than 370 million observations from hundreds of sites. This vast amount of data is currently under-utilised, with its main use being compliance monitoring, where a simple annual average is calculated and compared to the EU limits. More sophisticated analysis is hindered by the complexity of statistical analysis of air quality data, as will be discussed later. However, recent advances in the field of machine learning have resulting in an explosion of techniques for extracting patterns from large amounts of data, which have been applied in almost every field imaginable. These techniques offer the potential for extracting greater insight from the huge repository of air pollutant concentration data available for the UK.

1.4.1 Challenges associated with statistical analysis of ambient air quality data

Ambient air quality data has several unique characteristics that make its analysis challenging. Firstly, pollution events often result in spikes in pollutant concentrations in the time series that, while often acting as outliers with the potential to exert disproportionate leverage on the results, do not result from instrument errors, and should not be removed. As such, it is important to choose techniques that are robust to such outliers.

Autocorrelation is another characteristic of time series that should be considered, wherein error terms can be correlated over time, violating an assumption of many models, including the linear regression model, that the error terms be independent.

Additionally techniques must also be robust to missing data, which is a common problem in air quality data due to instrument faults etc.

However, the above problems are quite common in all time series data, and have been adequately solved elsewhere. The main difficulty with analysing a change in ambient concentration of an air pollutant, whether a long term trend or the effect of an intervention, is distinguishing that change from the background concentration and the effects of confounding factors. Meteorology almost invariably exerts the largest influence over pollutant concentrations, and variations in meteorology cause corresponding variations in concentration that can completely obscure the

much subtler long term trends. This frequently frustrates attempts to analyse the effects of interventions by comparing pollutant concentrations before and after its implementation, as was the case for the London Congestion Charge Zone (CSS), which was implemented in February 2003, unfortunately coinciding with a period of unusual meteorology in London that rendered any temporal comparison meaningless. (Beevers and Carslaw, 2005)

Other confounding factors, including temporal variation in the form of seasonal, weekly or diurnal cycles, long range transport, interactions with other pollutants, extreme events such as pollution episodes, concurrent changes in other emission sources/overlap in the implementation of other policy interventions, and changes in the instrumentation or methodology used to measure the pollutant add to the complexity of the time series. The nature and impact of these confounding factors are variable and depend on the location and pollutant under investigation. The first task of an analysis is often, therefore, disentangling the effects of various confounding factors, including meteorology and seasonality, on the time series in order to reveal the underlying trend.

Traditionally, methods such as calculating urban increments (Harrison et al., 2012; Kassomenos et al., 2014), applying seasonal trend decomposition (Carslaw, 2005; Bigi and Harrison, 2010), or training statistical models which include the confounder as predictors (Dijkema et al., 2008; Fensterer et al., 2014) have been applied to try and control for confounding variables.

However, modern statistical techniques in the field of machine learning offer the possibility for much greater accuracy and flexibility in accounting for confounders. This is because many machine learning models, such as ensemble trees like boosted regression trees and random forest, are non-linear, highly flexible and are intrinsically capable of including interactions between variables.

Boosted regression trees (BRT) were used by Carslaw and Taylor (2009) to infer the source characteristics of NO_x at Heathrow Airport, after controlling for meteorological factors, aircraft activity patterns and temporal factors by including them as variables in their models.

Grange et al. (2018) recently demonstrated that random forest models combined with bootstrapping could be used to ‘meteorologically normalise’ air pollutant concentration data (i.e. remove the influence of meteorological variables) to estimate the trend in concentrations under ‘average’ meteorological conditions.

These methods offer enormous potential for ameliorating one of the most difficult problems in air quality data analysis: that of distinguishing between changes in concentration driven by changes in emissions, and those driven by

other factors, such as meteorology. Chapters 3 and 4 explore this potential, as will be discussed later.

Finally, the data used in air quality data analysis is sourced from monitoring sites situated at particular locations. This data provides information about the local air quality at that location. The size of the area which the monitoring site can represent depends on the location: from just a single street for some urban road traffic sites, to a larger area for background sites. But it is sometimes useful to analyse air quality on a larger scale, such as over an entire network of sites, rather than a single monitoring site. For example, a detailed analysis of data from the London Marylebone site may provide a great deal of insight into changes in air quality on Marylebone Road, but what about air quality in London? In the UK? In Europe? Analysis of networks of sites offers the ability to answer these questions, and is discussed further in Sections 1.4.2 and 1.5.2, and in Chapter 2.

Within the field of air quality analysis, the following questions are of particular importance:

- How does air quality change over time?
- What effect do policy interventions, such as low emission zones, have on the concentrations of air pollutants?
- Is the level of air pollution in an area compliant with the legal limits? And if not, what factors are responsible for the non-compliance?

The methods for obtaining answers to each of these questions are examined in turn in the sections below.

1.4.2 Trend Analysis

Trend analysis involves examining the changes in pollutant concentration over time (Anttila and Tuovinen, 2010; Guerreiro et al., 2014). It can be used to track the effects of emission changes on the ambient air quality, and as evidence of the efficacy (or lack thereof) of policy interventions (Font and Fuller, 2016).

In its simplest aspect, the pollutant concentration as a function of time at a single monitoring site is analysed, either as a raw time series, or as a daily, monthly, or annual average.

However, the advantage of monitoring networks containing many - sometimes hundreds - of monitoring sites, is the ability to carry out trend analysis over a wider area or a longer time frame than may be represented by a single site.

Changes in air pollutant concentration at an individual monitoring site are affected by many factors in addition to emissions, such as meteorology, changes in the urban environment and dispersion. These confounding factors, as discussed in Section 1.4, impede an evaluation of the long term changes in air quality.

By including data from multiple sites over a wider area, the variation in concentration driven by site-specific factors cancel out, so that the changes driven by larger scale factors, such as changes in vehicle emissions or the results of interventions, can be seen.

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. This way, the changes in concentration common across many sites may be attributed to larger scale changes, while the changes only seen at a few sites, or sites geographically close together, might be interpreted as due to local confounders.

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 authors compared the trends at each site to establish a consensus and the differences between monitoring sites were rationalised using contextual information about each site. Other 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, with such a large number of comparables, it may be difficult to fully capture the patterns and anomalies (although cluster analysis has been used to ameliorate this problem (Malley et al., 2018)). It may be beneficial to aggregate data from multiple monitoring sites to gain a representative view of the average air quality.

Another method is to aggregate the data from multiple monitoring sites, then analyse the trend in the average concentration across many sites. Font and Fuller (2016) examined the trends in roadside increments of various pollutants between 2005-2009 and 2010-2014 by averaging data from 65 London monitoring sites. Data capture filters and linear interpolation were employed to ensure all individual time series from separate sites were of equal length.

The problem with this approach is that data filtering necessarily excludes some information from the analysis. In many monitoring networks, the time series at many sites may only span a few years, or sites may move so that time

series have little overlap. As a consequence, in areas with sparser monitoring site coverage, or for trend analysis of long time periods, there may not be many, or even any, sites that fulfill the data filtering requirements and therefore this approach may not be practical.

In such cases, it would be tempting to simply calculate the average concentration across all available monitoring sites to obtain a long term trend. However, if the monitoring sites have time series of various lengths and durations, the resulting trend 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.

In Chapter 2, a new method is presented that enables robust long term trends to be visualised using data from multiple monitoring sites with variable length time series. This rolling change method therefore ameliorates the disadvantages of both methods discussed previously: averaging the concentration across multiple sites reduces the impact of site specific confounders, while preserving more data in the analysis than data filtering allows. The method opens up the possibility of carrying out robust long term trend analysis for sparse monitoring networks with few, or no, long term monitoring sites.

The development of the rolling change method is timely given that low cost sensor networks are increasingly being used to monitor air quality. These sensors tend to be less reliable than traditional monitoring sites, as well as being more portable, and therefore such networks are even more vulnerable to the problems associated with site flux. As such, the rolling change method could be valuable in analysis of trends using data from these networks.

Moreover, the issue of aggregating data from time series of differing lengths is not exclusive to the analysis of air quality data. The method could find application in any discipline dealing with multiple time series of differing lengths.

To facilitate the use of the rolling change method, an open source R package named `aqtrends` was developed. The package contains a function that outputs the rolling trend for a given data set input. More details about `aqtrends` is given in Appendix III.

1.4.3 Intervention Analysis

1.4.3.1 Air quality interventions

In response to dangerously high levels of air pollution in urban areas across the world, governments have implemented a range of interventions aimed at controlling emissions.

Many of these interventions focus on reducing emissions from road traffic. At the national level, increasingly stringent limits on the emissions of major pollutants, such as particulate matter and NO₂, have been placed on vehicle manufacturers. For instance, the EU has imposed a series of emission standards, from Euro 1 to Euro 6, in order to encourage technological innovation and implementation in reducing emissions from manufacturers.

At the other end of the scale, interventions have also been applied within individual cities and highways to control emissions in particularly polluted areas. These interventions include low emission zones (LEZ) and clean air zones (CAZ), variable speed limits, and reduction in speed limits, as well as other mechanisms for reducing congestion and stop-start traffic flows.

Low emission zones and clean air zones are areas where vehicles not meeting certain emission standards are prohibited from entering, and, failing this, are subject to fines. As of 2019, there were about 250 active LEZs across Europe, including in London, Paris, Brussels, Lisbon, Madrid, Barcelona, Berlin, Munich, Amsterdam, Oslo, Stockholm and Athens. The LEZs vary widely in size, the severity of the restriction and whether they are permanent, seasonal or emergency (activated only during periods of high pollution) (McGrath, 2019).

The London LEZ is an intervention aimed at reducing air pollution in London. Figure 1.2 shows the area covered by the London LEZ. Its staged implementation began in February 2008 with Phase 1, which mandated that heavy goods vehicles (HGVs) over 12 tonnes in weight must meet the Euro III emission standard. Phase 2, which began in July 2008, extended this to all lorries over 3.5 tonnes, as well as buses and coaches. Phase 3 activated in October 2010, and specified that larger light goods vehicles (LGVs) and minibuses would also be covered by the scheme. The final Phase 4 began in January 2012, and tightened the minimum standard to Euro IV for PM for lorries over 3.5 tonnes, buses and coaches (TfL, 2008).

1.4.3.2 Analysing the impact of interventions

Studies investigating the impacts of policy interventions on air quality (and, sometimes, human health) are called 'accountability studies'. These studies

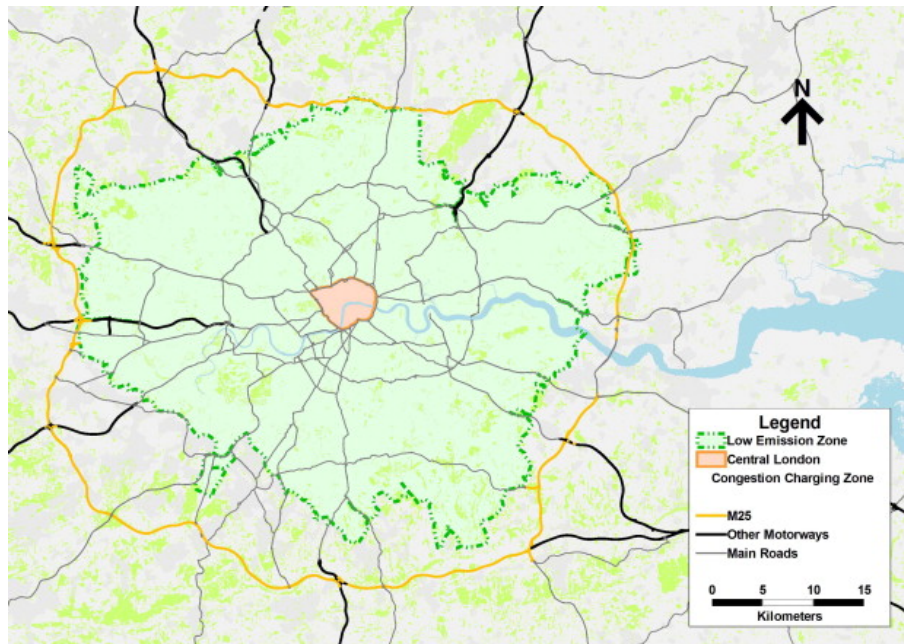


Figure 1.2: Extent of the London LEZ. Taken from Ellison et al. (2013)

formulate the problem as a chain, and trace the changes from the initial regulation to the emissions of air pollutants, to the ambient air quality, to the exposure, to the health impacts, as shown in Figure 1.3 (Henneman et al., 2017). An accountability study may investigate part or all of the chain.

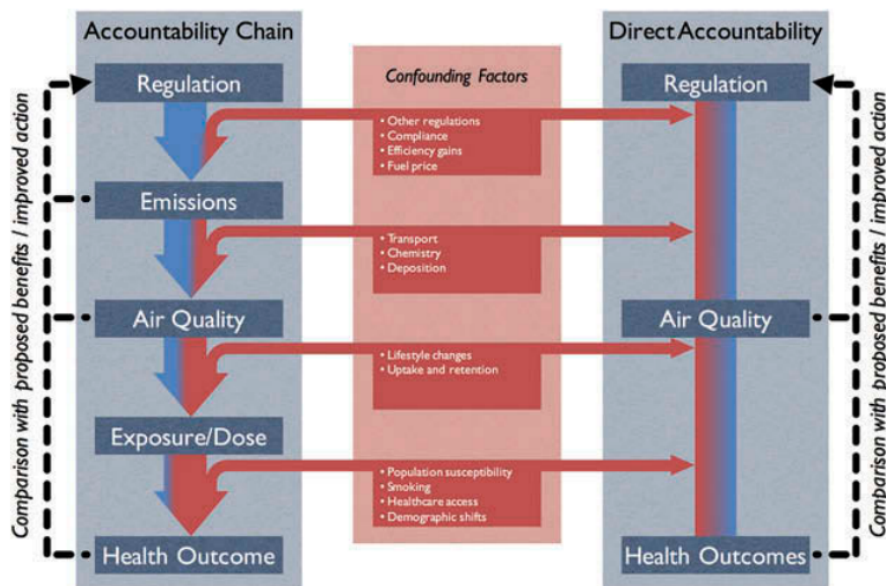


Figure 1.3: The accountability chain (Henneman et al., 2017). Confounding factors that may affect each link of the chain are shown in red. This diagram is from the paper by Henneman et al. (2017).

Analysis and quantification of the response of a target variable (e.g. ambient concentration of an air pollutant) to changes in a causal driver (e.g. implementation of a policy intervention) can be ascertained by comparison with an appropriate control scenario (Greenbaum and Org, 2017; Henneman et al., 2017). In order to establish what changes have been brought about by an intervention, there must be information on what would have happened if the intervention had not been in place to compare it to. In scientific laboratory experiments, the experiment is often performed twice, with and without the intervention, but with all other conditions kept the same. In this way, the effect of the intervention can be distinguished from the effects of changes in other factors (confounders). The situation is more complex for studies of air quality interventions: it is impossible, for example, to measure the pollutant concentrations that would have occurred had a low emission zone not been in force if, indeed, a low emission zone has been implemented at a particular site. Consequently, such studies have historically made use of two methods for estimating a counterfactual (or control) scenario. A temporal control uses the period of time before the intervention was implemented as the control scenario. Alternatively, the air quality at a similar location where the intervention was not implemented can provide a spatial control scenario.

Unfortunately, both temporal and spatial controls are vulnerable to the effects of confounders.

One of the main challenges associated with air quality research, including accountability studies, is distinguishing the effects of causal drivers of changes in air quality (e.g. an intervention) from confounding variables which may also change over the period of analysis. Each link in the accountability chain is vulnerable to confounding by variation in factors that impact air quality. For example, the relationship between the regulation and emissions may be confounded by other concurrent factors that affected vehicle use over the same period, such as changes in fuel price or the implementation of other traffic policies, or by changes in vehicle emission technologies. The relationship between emissions and ambient air quality is confounded by seasonal variation, local meteorology, long range transport and other atmospheric processes (Henneman et al., 2017).

A comparison of the data with temporal and spatial controls will be confounded respectively by temporal and spatial variation of the confounding variables. For example, changes in ambient air quality before and after an intervention was implemented will be driven not only by the effect of the intervention, but also by variation in meteorology, long-range transport, road-user behaviour, emission

technologies, and any other variables impacting on air quality that vary over the time period. Similarly, comparison with a spatial control will be influenced by confounders that vary spatially between the two locations: for example, local meteorology and dispersion, or differences in the urban environment, demographics and road-user behaviour.

The effect of the confounders on air quality are difficult to distinguish from the effect of the intervention, and therefore it is challenging to attribute changes in air quality to any single driver. Any attempt to analyse or quantify the impact of a single variable, such as the presence of the intervention, on the ambient air quality must find a way of removing the influence of such confounders, otherwise there is no way to determine whether the observed changes are driven by the presence of the intervention, or by variation of one or more of the confounding variables.

One method for removing the effect of confounders is by modelling the ambient concentration as a function of the confounders and the intervention (often represented by an indicator variable, that takes the value 0 when the intervention is inactive, and 1 when it is active). The modelled relationship between the intervention and the ambient concentration can then be analysed, for example, in a linear model, the effect of the intervention is estimated by the coefficient of the intervention variable (Dijkema et al., 2008; Fensterer et al., 2014; Wolff, 2014).

Dijkema et al. (2008) analysed the effect of a speed limit intervention on the Amsterdam ring highway on NO_x , PM and black carbon concentrations using a linear model with the predictor variables traffic flow, traffic congestion, wind direction and an indicator variable representing the intervention (whether or not it was in effect). The authors' estimation of the effect of the intervention was represented by the coefficient of the indicator variable. (Dijkema et al., 2008)

The study by Fensterer et al. (2014) demonstrated an interesting extension of the standard linear model, where PM_{10} concentrations in Munich were modelled using semi-parametric regression. Meteorological conditions, background concentrations, public holidays, wind direction and diurnal variation were used as input variables in a linear model along with the indicator, but the coefficients of the wind direction and diurnal variation were non-linear non-parametric functions (P-splines and cyclic P-splines, respectively). This enabled highly non-linear relationships to be modelled with greater accuracy and flexibility.

However, a weakness of these studies was that the relationships between ambient air pollutant concentrations and the confounding variables are often complex, interdependent and highly non-linear. Even the use of non-linear functions as

coefficients in the model, as in Fensterer et al. (2014), make assumptions about the relationships between variables that may not be true, and may vary widely across different locations and pollutants. Furthermore, air quality data rarely satisfied the assumptions of the linear model, being interdependent, non-normal (a long tail is common due to the presence of extreme data points caused by pollution events), heteroscedastic and autocorrelated. For these reasons, the most commonly used method for modelling these systems, multivariate linear regression, is not ideal. The accuracy of the models, and therefore the ability to make causal inferences about ambient air quality, can be improved by the use of more sophisticated models that are robust to the properties of air quality monitoring data, and capable of incorporating non-linearities and interactions.

Ensemble decision tree methods, such as random forest models, have been found to perform well in air quality domains (Grange et al., 2018; Suleiman et al., 2016). The advantages of these ensemble decision trees make them particularly well suited to air quality applications. The models are capable of modelling non-linear relationships and interactions between the variables which are common in air quality data. Additionally, any type of predictor variable can be used as inputs. The model is unaffected by transformations or differences in scale among the predictors and automatically ignores irrelevant predictors, simplifying the pre-treatment of data prior to modelling as scaling and feature reducing are unnecessary. Decision trees can be used to model step changes, are robust to missing data and outliers and automatically model interactions between variables. Crucially, they are among the most interpretable of all machine learning models, offering the ability to infer a great deal of information from the fitted model as well as making accurate predictions. (Carslaw and Taylor, 2009).

Carslaw et al. (2012) used boosted regression trees to quantify the impact of the closure of Heathrow airport following the eruption of Eyjafjallajökull in 2010 on NO_x and NO_2 concentrations at London sites. A boosted regression tree model of NO_x and NO_2 concentration was built using pre-airport closure data, then used to model NO_x and NO_2 concentrations under a business-as-usual scenario during the closure, and compared to the actual measured concentrations over that time period. From this, the authors were able to detect a significant change in concentration over the short time period for which the airport was closed at some sites, and noted the potential for the method's application to other, longer term interventions. (Carslaw et al., 2012)

In another study, Lana et al. (2016) used random forest models to measure the relative contributions of traffic emissions, meteorological factors, and background

concentrations to air quality in Madrid, in order to establish the potential efficacy of proposed traffic restrictions on air quality. The relative importance of input variables was quantified by fitting the model with different combinations of input features and comparing the error rates. The authors concluded that, at most sites, traffic emissions were not very significant as a source relative to meteorology and had very little impact on the regional background, leading to the recommendation of permanent low emission zones as an alternative to traffic restrictions as an air quality intervention. (Lana et al., 2016)

Despite the urgency of the issue, the efficacy of interventions such as low emission zones on the ambient air quality remains an open question. Numerous studies have reached different conclusions on the matter (Malina and Scheffler, 2015; Ferreira et al., 2015; Cyrus et al., 2014; Boogaard et al., 2012; Panteliadis et al., 2014). While the unique circumstances of the studied area and the robustness of the policy's implementation are important factors, one reason for this ambiguity is the difficulty of such analyses, as discussed previously. Both the impact of confounding variables on the air pollutant concentration, and the challenge of establishing a control scenario hinder the ability of these studies to reach definitive conclusions.

In Chapter 3, an intervention analysis of the London LEZ is presented. The method uses non-linear random forest models to model the PM₁₀ concentration to account for the confounding variables. An indicator variable is included representing the stages of the LEZ, and used to estimate 'business as usual' control scenarios for comparison with the meteorologically normalised real data.

1.4.4 Compliance Monitoring

The most common use of ambient air quality monitoring data is for monitoring compliance with the European Union's air quality limits. The EU imposes limits on the concentrations of various air pollutants and requires all member states to annually report the annual (and sometimes daily) average concentrations. Imposing limits on the concentrations of air pollutants incentivises governments to prioritise policies aimed at improving air quality.

The limit values are based on the World Health Organisation (WHO) guidelines which estimate the maximum level of pollution above which adverse health effects may be experienced. Given the scale and complexity of the problem of reducing emissions, many of the EU limit values are set higher than the WHO recommendation in order to set more realistic targets. For example, based on its most recent evaluation of the scientific evidence, WHO recommend that the

annual mean concentrations of PM_{2.5} and PM₁₀ not exceed 10 $\mu\text{g m}^{-3}$ and 20 $\mu\text{g m}^{-3}$ respectively, and that the O₃ 8-hour mean not exceed 100 $\mu\text{g m}^{-3}$, all of which are lower than the EU limits shown in Table 1.2 (WHO, 2006; European Commission, 2019). Despite this, exceedences of the limit values are frequently reported in many member states. In 2018, six countries, including the UK, were referred to the European Court of Justice for consistent failure to meet the air quality limits (Carrington, 2018).

Table 1.2 shows the limit values and metrics of each air pollutant legislated by the EU. It can be seen that while for some species, such as O₃, the concentration is reported as the daily maximum of the running 8-hour mean concentration¹, for others, such as NO₂, only a simple annual average is required (European Commission, 2019; AQEG, 2009)..

Table 1.2: European Directive limit values of the concentration of selected air pollutants (European Commission, 2019)

Pollutant	Concentration ($\mu\text{g m}^{-3}$)	Averaging period	Permitted yearly ex- ceedences
Fine particles (PM _{2.5})	25	1 year	n/a
SO ₂	350	1 hour	24
SO ₂	125	24 hours	3
NO ₂	200	1 hour	18
NO ₂	40	1 year	n/a
PM ₁₀	50	24 hours	35
PM ₁₀	40	1 year	n/a
Benzene	5	1 year	n/a
O ₃	120	Maximum daily 8 hour mean	25 days averaged over 3 years

As discussed previously, the ambient concentrations of air pollutants are influenced by many confounding factors other than emissions, in particular, meteorology. Since meteorology varies from year to year, some variation in the annual average concentration of air pollutants will be due to these interannual variations in meteorology, rather than variations in emissions. This means that changes in the annual average concentration cannot be directly attributed to changes in emissions. More importantly, compliance with the EU limits could depend on the meteorology of that year. For example, during a year with meteorology that drives high concentrations of air pollutants - such as a particularly cold winter characterised by prolonged periods of low dispersion

¹The running 8-hour mean is assigned the date of the last hour of the running mean

conditions - the reported annual average concentrations of primary pollutants such as NO_2 may be very high even if the emissions of this species are not unusually high. In such cases, exceedence of the air quality limits may be a consequence of meteorology rather than emissions.

It can therefore be seen that the issue of the extent to which interannual meteorological variation affects the annual average concentration of air pollutants is very important, but it is largely ignored. The difficulty of separating the influence of meteorology (or other confounding factors) from the influence of emissions when analysing ambient monitoring data has been discussed previously. For this reason, quantitatively determining the variation in average concentration due to meteorological variation is challenging. Chapter 4 aims to address this by quantifying that variation, and therefore evaluating the degree to which compliance is affected by meteorology.

In Chapter 4, a new method is presented to quantitatively estimate the variation in the ambient concentration due to meteorological variation, involving a novel application of the random forest meteorological normalisation method described in Chapter 3. Novel visualisation tools using heatmaps and cumulative sum plots are shown to enable detailed investigation of which years were characterised by 'good' meteorology (that which drove better-than-usual air quality) or 'bad' meteorology, as well as the underlying reasons for this. Based on this quantification, the effect of meteorological variation on site compliance or exceedence in a given year can be estimated.

1.5 Structure and Content of this Thesis

1.5.1 Objectives

The purpose of this thesis is to develop new methods for analysing air pollutant concentration data, using modern machine learning techniques to overcome some of the traditional problems associated with such analyses. Subsequently, these new methods are applied to the vast amount of ambient air quality data routinely collected in the UK to derive new insights about changes in air quality in the UK.

Chapter 2 describes a new technique that enables data from multiple monitoring sites in a network to be used in trend analysis, even in the presence of bias due to site flux. The efficacy of the technique is demonstrated through a case study of air quality trends at roadside in London, and applied to analyse trends in roadside air quality in Scotland and the UK.

Chapters 3 and 4 attempt to extract more information from ambient air quality monitoring data by using random forest to remove the influence of confounding factors such as meteorology. Chapter 3 uses this meteorological normalisation technique alongside a simulated counterfactual scenario (as demonstrated in [Carslaw et al. \(2012\)](#)) to attempt to evaluate the efficacy of the London LEZ in reducing air pollution. Chapter 4 combines the random forest meteorological normalisation method with bootstrapping techniques and novel visualisation methods to calculate the interannual meteorological variation and evaluate its effect on the annual average concentration of air pollutants.

Sophisticated statistical analysis of air quality data enables rigorous trend analysis that can be used to draw conclusions about the drivers of trends. In Chapters 2 and 3, changes in air quality are compared with the implementation of policies aimed at reducing air pollution, such as low emission zones and vehicle emission technologies, and used to make cautious inferences about the efficacy of such policies.

Chapter 4 also applies the new methods to critically evaluate current methodologies for compliance monitoring of air quality. Currently, ambient air quality data is used to calculate annual averages (or rolling averages for some pollutants e.g. O₃) which are compared with a numeric limit. In Chapter 4, the disadvantage of this approach is demonstrated through an analysis of the impact of interannual meteorological variation on the annual average concentration of common air pollutants.

The following sections provide a brief description of each chapter, with a more detailed introduction given at the beginnings of the chapters themselves.

1.5.2 A new method for trend analysis in biased monitoring networks

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](#)). While the data from a single monitoring station provides information about the local variation in air quality at a specific location, aggregation of data from multiple monitoring stations enables the effects of local variability to be ‘averaged out’, leaving a better indication of the large-scale trend.

However, simply calculating the average concentration across every monitoring site can result in a biased trend that does not represent the changes in air pollutant

concentration at the individual monitoring sites due to biases in the monitoring network caused by site flux (movement of monitoring sites) during the period of analysis. In Chapter 2, this biasing effect is demonstrated using simulated data, and new techniques to identify and mitigate the biasing effect of variation in time series length on the trend in average concentration are developed. The efficacy of the method is demonstrated using simulated data and a trend analysis of NO_x and NO_2 concentrations in London between 2000-2017 using data from the entire London monitoring network.

1.5.3 Robust analysis of trends and policy interventions using random forest

Governments are under pressure from both legislative bodies and the public to improve air quality in urban areas. The policies used to do this usually focus on traffic control, such as establishing low emission zones, clean air zones, variable speed limits and reduction of speed limits.

The potential efficacy of these interventions are evaluated using atmospheric modelling during the planning stages, but there is a paucity of methods capable of analysing their actual effect on ambient air pollution after their introduction.

One such method is presented in Chapter 3. The method uses the ‘meteorological normalisation’ method developed by [Grange et al. \(2018\)](#) with an indicator variable for the presence of the intervention to generate pollutant concentration trends for different counterfactual scenarios, which can be compared.

1.5.4 Quantification of the effect of interannual meteorological variation on air pollutant concentration, and the implications for compliance metrics

Meteorological variation has a huge influence on air pollutant concentration. There is therefore a concern that variation in meteorology from year-to-year could influence the annual average pollutant concentration: the metric that is used to measure compliance with limits for many air pollutants. If this is the case, then there is a risk that compliance or exceedance with the limit values could be driven by interannual meteorological variation, rather than an increase in emissions.

Chapter 4 investigates this possibility by attempting to quantify the effect of interannual meteorological on concentrations of air pollutants at monitoring sites in the UK between 2008 and 2017. Techniques involving the use of heatmaps and

CUSUM plots are developed for this purpose, and the results of the analysis are used to assess whether variations in meteorology are responsible for concentration limit exceedances from year to year.

Bibliography

- Adams, K., Greenbaum, D. S., Shaikh, R., van Erp, A. M., Russell, A. G., 2015. Particulate matter components, sources, and health: Systematic approaches to testing effects. *Journal of the Air & Waste Management Association* 65 (5).
URL <https://doi-org.libproxy.york.ac.uk/10.1080/10962247.2014.1001884>
- Airqualitynews.com, 2018. UK government to face EU legal action over air pollution.
URL <https://airqualitynews.com/2018/05/17/uk-government-to-face-eu-legal-action-over-air-pollution/>
- Anderson, J. O., Thundiyil, J. G., Stolbach, A., 2012. Clearing the Air: A Review of the Effects of Particulate Matter Air Pollution on Human Health. *Journal of Medical Toxicology* 8 (2), 166–175.
- Anttila, P., Tuovinen, J.-P., 2010. Trends of primary and secondary pollutant concentrations in Finland in 1994–2007. *Atmospheric Environment* 44 (1), 30–41.
URL <https://doi.org/10.1016/j.atmosenv.2009.09.041>
- AQEG, 2004. Nitrogen dioxide in the United Kingdom. Tech. rep., Air Quality Expert Group.
URL <https://uk-air.defra.gov.uk/library/aqeg/publications>
- AQEG, 2005. Particulate matter in the United Kingdom. Tech. rep., Air Quality Expert Group.
URL <https://uk-air.defra.gov.uk/library/aqeg/publications>
- AQEG, 2009. Ozone in the United Kingdom. Tech. rep., Air Quality Expert Group.
URL <https://uk-air.defra.gov.uk/library/aqeg/publications>
- AQEG, 2012. Fine Particulate Matter (PM_{2.5}) in the United Kingdom. Tech. rep., Air Quality Expert Group.
URL <https://uk-air.defra.gov.uk/library/aqeg/publications>

- AQEG, 2019. Non-Exhaust Emissions from Road Traffic. Tech. rep., Air Quality Expert Group.
URL https://uk-air.defra.gov.uk/library/reports.php?report_{_}id=992
- AQEG, 2020. Non-methane volatile organic compounds in the UK. Tech. rep., Air Quality Expert Group.
URL <https://uk-air.defra.gov.uk/library/aqeg/publications>
- Beevers, S. D., Carslaw, D. C., 2005. The impact of congestion charging on vehicle emissions in London. *Atmospheric Environment* 39 (1), 1–5.
URL <https://doi.org/10.1016/j.atmosenv.2004.10.001>
- Bigi, A., Harrison, R. M., 2010. Analysis of the air pollution climate at a central urban background site. *Atmospheric Environment* 44 (16), 2004–2012.
URL <https://doi.org/10.1016/j.atmosenv.2010.02.028>
- Boogaard, H., Janssen, N. A., Fischer, P. H., Kos, G. P., Weijers, E. P., Cassee, F. R., van der Zee, S. C., de Hartog, J. J., Meliefste, K., Wang, M., Brunekreef, B., Hoek, G., 2012. Impact of low emission zones and local traffic policies on ambient air pollution concentrations. *Science of The Total Environment* 435-436, 132–140.
URL <https://doi.org/10.1016/j.scitotenv.2012.06.089>
- Carrington, D., 2018. UK taken to Europe’s highest court over air pollution.
URL <https://www.theguardian.com/environment/2018/may/17/uk-taken-to-europes-highest-court-over-air-pollution>
- Carslaw, D. C., 2005. On the changing seasonal cycles and trends of ozone at Mace Head, Ireland. *Atmospheric Chemistry and Physics* 5, 3441–3450.
URL www.atmos-chem-phys.org/acp/5/3441/
- Carslaw, D. C., Murrell, T. P., Andersson, J., Keenanc, M., 2016. Have vehicle emissions of primary NO₂ peaked? *Faraday Discussions* 189, 439–454.
- Carslaw, D. C., Taylor, P. J., 2009. Analysis of air pollution data at a mixed source location using boosted regression trees. *Atmospheric Environment* 43 (22-23), 3563–3570.
URL <https://doi.org/10.1016/j.atmosenv.2009.04.001>
- Carslaw, D. C., Williams, M. L., Barratt, B., 2012. A short-term intervention study - Impact of airport closure due to the eruption of Eyjafjallajökull on near-field air

- quality. *Atmospheric Environment* 54, 328–336.
URL <https://doi.org/10.1016/j.atmosenv.2012.02.020>
- Cyrys, J., Peters, A., Soentgen, J., Wichmann, H.-E., 2014. Low emission zones reduce PM₁₀ mass concentrations and diesel soot in German cities. *Journal of the Air & Waste Management Association* 64 (4), 481–487.
URL <https://www.tandfonline.com/doi/full/10.1080/10962247.2013.868380>
- Defra, 2019. UK and EU Air Quality Limits.
URL <https://uk-air.defra.gov.uk/air-pollution/uk-eu-limits>
- Dijkema, M. B. A., Van Der Zee, S. C., Brunekreef, B., Van Strien, R. T., 2008. Air quality effects of an urban highway speed limit reduction. *Atmospheric Environment* 42, 9098–9105.
URL <https://doi.org/10.1016/j.atmosenv.2008.09.039>
- Ellison, R. B., Greaves, S. P., Hensher, D. A., 2013. Five years of London's low emission zone: Effects on vehicle fleet composition and air quality. *Transportation Research Part D: Transport and Environment* 23, 25–33.
URL <https://doi.org/10.1016/j.trd.2013.03.010>
- European Commission, 2019. Air quality standards.
URL <https://ec.europa.eu/environment/air/quality/standards.htm>
- European Union, 2008. Directive 2008/50/EC of the European Parliament and of the Council of 21 May 2008 on ambient air quality and cleaner air for Europe.
URL <https://eur-lex.europa.eu/eli/dir/2008/50/2015-09-18>
- Fensterer, V., Küchenhoff, H., Maier, V., Wichmann, H.-E., Breitner, S., Peters, A., Gu, J., Cyrys, J., 2014. Evaluation of the impact of low emission zone and heavy traffic ban in munich (germany) on the reduction of pm₁₀ in ambient air. *International Journal of Environmental Research and Public Health* 11 (5), 5094–5112.
URL <http://www.mdpi.com/1660-4601/11/5/5094/>
- Ferreira, F., Gomes, P., Tente, H., Carvalho, A., Pereira, P., Monjardino, J., 2015. Air quality improvements following implementation of Lisbon's Low Emission Zone. *Atmospheric Environment* 122, 373–381.
URL <https://doi.org/10.1016/j.atmosenv.2015.09.064>

- Font, A., Fuller, G. W., 2016. Did policies to abate atmospheric emissions from traffic have a positive effect in London? *Environmental Pollution* 218, 463–474.
URL <https://doi.org/10.1016/j.envpol.2016.07.026>
- Grange, S. K., Carslaw, D. C., Lewis, A. C., Boleti, E., Hueglin, C., 2018. Random forest meteorological normalisation models for Swiss PM₁₀ trend analysis. *Atmospheric Chemistry and Physics Discussions*, 1–28.
URL <https://doi.org/10.5194/acp-18-6223-2018>
- Greenbaum, D. S., Org, D., 2017. Learning about “cause” and “effect” through well-designed studies of air quality interventions. *International Journal of Public Health* 62, 719–720.
URL <https://link.springer.com/article/10.1007/s00038-017-0979-0>
- Guerreiro, C. B., Foltescu, V., de Leeuw, F., 2014. Air quality status and trends in Europe. *Atmospheric Environment* 98, 376–384.
URL <https://doi.org/10.1016/j.atmosenv.2014.09.017>
- Harrison, R. M., Laxen, D., Moorcroft, S., Laxen, K., 2012. Processes affecting concentrations of fine particulate matter (PM_{2.5}) in the UK atmosphere. *Atmospheric Environment* 46, 115–124.
URL <https://doi.org/10.1016/j.atmosenv.2011.10.028>
- Henneman, L. R., Liu, C., Mulholland, J. A., Russell, A. G., 2017. Evaluating the effectiveness of air quality regulations: A review of accountability studies and frameworks. *Journal of the Air and Waste Management Association* 67 (2), 144–172.
URL <http://www.tandfonline.com/action/journalInformation?journalCode=uawm20>
<http://www.tandfonline.com/loi/uawm20>
<http://dx.doi.org/10.1080/10962247.2016.1242518>
- Kassomenos, P. A., Vardoulakis, S., Chaloulakou, A., Paschalidou, A. K., Grivas, G., Borge, R., Lumbreras, J., 2014. Study of PM₁₀ and PM_{2.5} levels in three European cities: Analysis of intra and inter urban variations. *Atmospheric Environment* 87, 153–163.
URL <https://doi.org/10.1016/j.atmosenv.2014.01.004>
- Lana, I., Del Ser, J., Padro, A., Velez, M. V., Casanova-Mateo, C., 2016. The role of local urban traffic and meteorological conditions in air pollution: A data-based case study in Madrid, Spain. *Atmospheric Environment* 145, 424–438.
URL <https://doi.org/10.1016/j.atmosenv.2016.09.052>

- Malina, C., Scheffler, F., 2015. The impact of Low Emission Zones on particulate matter concentration and public health. *Transportation Research Part A: Policy and Practice* 77, 372–385.
URL <https://doi.org/10.1016/j.tra.2015.04.029>
- 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. *Atmospheric Chemistry and Physics* 18, 3563–3587.
URL <https://doi.org/10.5194/acp-18-3563-2018>
- Masiol, M., Squizzato, S., Formenton, G., Harrison, R. M., Agostinelli, C., 2017. Air quality across a European hotspot: Spatial gradients, seasonality, diurnal cycles and trends in the Veneto region, NE Italy. *Science of The Total Environment* 576, 210–224.
URL <https://doi.org/10.1016/j.scitotenv.2016.10.042>
- Mavroidis, I., Chaloulakou, A., 2011. Long-term trends of primary and secondary NO₂ production in the Athens area. Variation of the NO₂/NO_x ratio. *Atmospheric Environment* 45 (38), 6872–6879.
URL <https://doi.org/10.1016/j.atmosenv.2010.11.006>
- McGrath, M., 2019. ULEZ: How does London's new emissions zone compare?
URL <https://www.bbc.co.uk/news/science-environment-47816360>
- Monks, P. S., 2005. Gas-phase radical chemistry in the troposphere. *Chemical Society Reviews* 34 (5), 376–395.
- NAEI, 2020. UK National Atmospheric Emissions Inventory.
URL <http://naei.beis.gov.uk/index>
- Panteliadis, P., Strak, M., Hoek, G., Weijers, E., van der Zee, S., Dijkema, M., 2014. Implementation of a low emission zone and evaluation of effects on air quality by long-term monitoring. *Atmospheric Environment* 86, 113–119.
URL <https://doi.org/10.1016/j.atmosenv.2013.12.035>
- PORG, 1997. Ozone in the United Kingdom. Fourth Report of the UK Photochemical Oxidants Review Group. Tech. rep., Department of the Environment, Transport and the Regions, London.
URL <http://cedadocs.ceda.ac.uk/1100/>

BIBLIOGRAPHY

Royal College of Physicians, 2016. Every breath we take: the lifelong impact of air pollution. Tech. rep., RCP.

URL <https://www.rcplondon.ac.uk/projects/outputs/every-breath-we-take-lifelong-impact-air-pollution>

Schultz, M. G., Akimoto, H., Bottenheim, J., Buchmann, B., Galbally, I. E., Gilge, S., Helmig, D., Koide, H., Lewis, A. C., Novelli, P. C., Plass-Dülmer, C., Ryerson, T. B., Steinbacher, M., Steinbrecher, R., Tarasova, O., Tørseth, K., Thouret, V., Zellweger, C., 2015. The Global Atmosphere Watch reactive gases measurement network. *Elementa Science of the Anthropocene* 3.

URL <https://doi.org/10.12952/journal.elementa.000067>

Suleiman, A., Tight, M. R., Quinn, A. D., 2016. Hybrid Neural Networks and Boosted Regression Tree Models for Predicting Roadside Particulate Matter. *Environmental Model Assessment* 21, 731–750.

URL <https://link.springer.com/content/pdf/10.1007/s10666-016-9507-5.pdf>

TfL, 2008. London Low Emission Zone: Impacts Monitoring Baseline Report. Tech. rep., Transport for London.

URL <http://content.tfl.gov.uk/lez-impacts-monitoring-baseline-report-2008-07.pdf>

The World Bank and Institute for Health Metrics and Evaluation, 2016. The Cost of Air Pollution: Strengthening the Economic Case for Action. Tech. rep., University of Washington, Seattle.

URL <http://documents.worldbank.org/curated/en/781521473177013155/The-cost-of-air-pollution-strengthening-the-economic-case-for-action>

Wang, P., Chen, Y., Hu, J., Zhang, H., Ying, Q., 2019. Attribution of Tropospheric Ozone to NO_x and VOC Emissions: Considering Ozone Formation in the Transition Regime. *Environmental Science & Technology* 53, 1404–1412.

WHO, 2006. WHO Air quality guidelines for particulate matter, ozone, nitrogen dioxide and sulfur dioxide. Tech. rep., World Health Organisation.

URL <https://www.euro.who.int/en/health-topics/environment-and-health/air-quality/activities/update-of-who-global-air-quality-guidelines>

Wolff, H., 2014. Keep Your Clunker in the Suburb: Low-emission Zones and

BIBLIOGRAPHY

Adoption of Green Vehicles. *The Economic Journal* 124 (578), F481–F512.
URL <http://doi.wiley.com/10.1111/eoj.12091>

2. A new trend analysis approach for air quality network data

2.1 Introduction

2.1.1 Background

Air quality monitoring networks are instrumental in 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 means 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. 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 (the average trend) over monitoring sites of differing duration is sensitive to biases in the monitoring network. Air pollution monitoring sites are frequently 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*¹ 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 NO₂ and PM₁₀ 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.

¹Defined as the change in number of sites over time.

In this chapter, 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 2000-2017 using data from roadside monitoring sites in the London air quality network.

London was chosen as a 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.

The methods were then applied to an analysis of the trends in roadside air quality on a larger scale. The long term trends in NO_x , NO_2 and PM_{10} roadside concentration, and in the NO_2/NO_x ratio at the roadside in Scotland and across the entire UK between 2000 and 2017 were analysed. This monitoring network has experienced considerable growth over this period, therefore the rolling change method is an appropriate choice for the analysis.

2.2 Method

2.2.1 Identification of bias effects on the trend

Evidence of a bias in trends from 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.

2.2.2 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 continuously 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 considerable amount of 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 2.2.1. 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 2000–2017 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 continuously 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 2015–2017 to calculate the concentration change in 2016.

Similarly to the data filtering method described in Section 2.1.1, 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.

Figure 2.1 shows a schematic of the process, while Equations 2.1 and 2.2 describe the rolling regression and the recursive concentration calculation respectively. A more detailed algorithm can be found in Appendix I.

The terms in Equations 2.1 and 2.2 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 2.3 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).

$$Y_i = \beta_0 + \beta_i X_i + \epsilon_i \quad (2.1)$$

$$\Delta y_i = \Delta y_{i-1} + \beta_i \quad (2.2)$$

$$\Delta y_i = f(x_i) + \epsilon_i \quad (2.3)$$

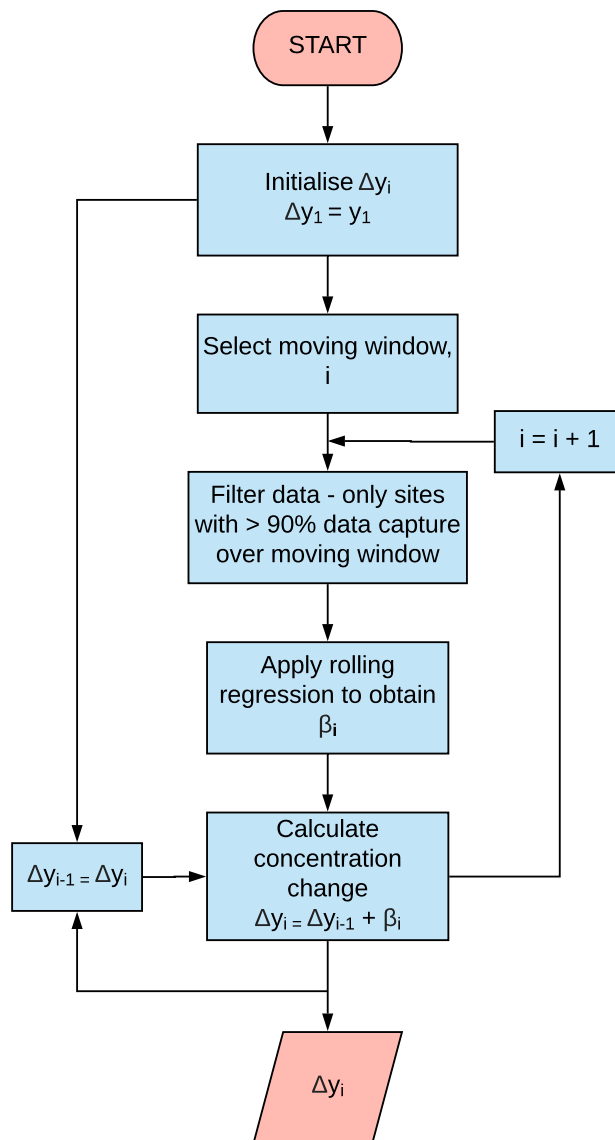


Figure 2.1: Schematic of the rolling change method. The output for the process as a whole (the concentration change for the rolling window, i ,) is highlighted in red.

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).

2.2.3 Description of data

The data used in both the London and the UK case studies 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 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 London trend analysis. The mean and the standard deviation of the hourly NO_x and NO_2 concentrations measured at each London monitoring site is given in Appendix II (Table B.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 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. 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 Appendix II (Table B.1).

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

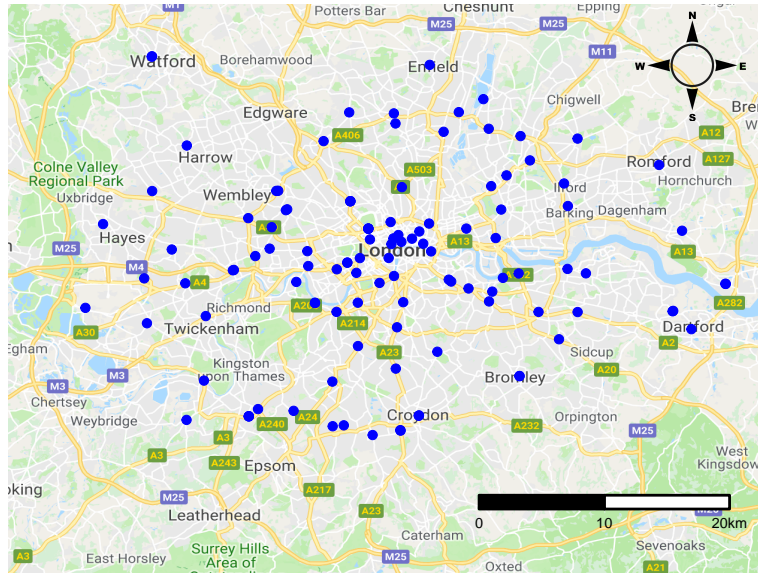


Figure 2.2: Map showing the locations of the London roadside monitoring sites measuring NO_x and NO_2 used in the analysis. More information on individual sites can be found in Appendix II (Table B.1).

2.3 Results and Discussion

2.3.1 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 2000-2015. 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 $y_i = 1 + 0.08x_i + \epsilon_i$ 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, 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 Figure 2.3.

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%.

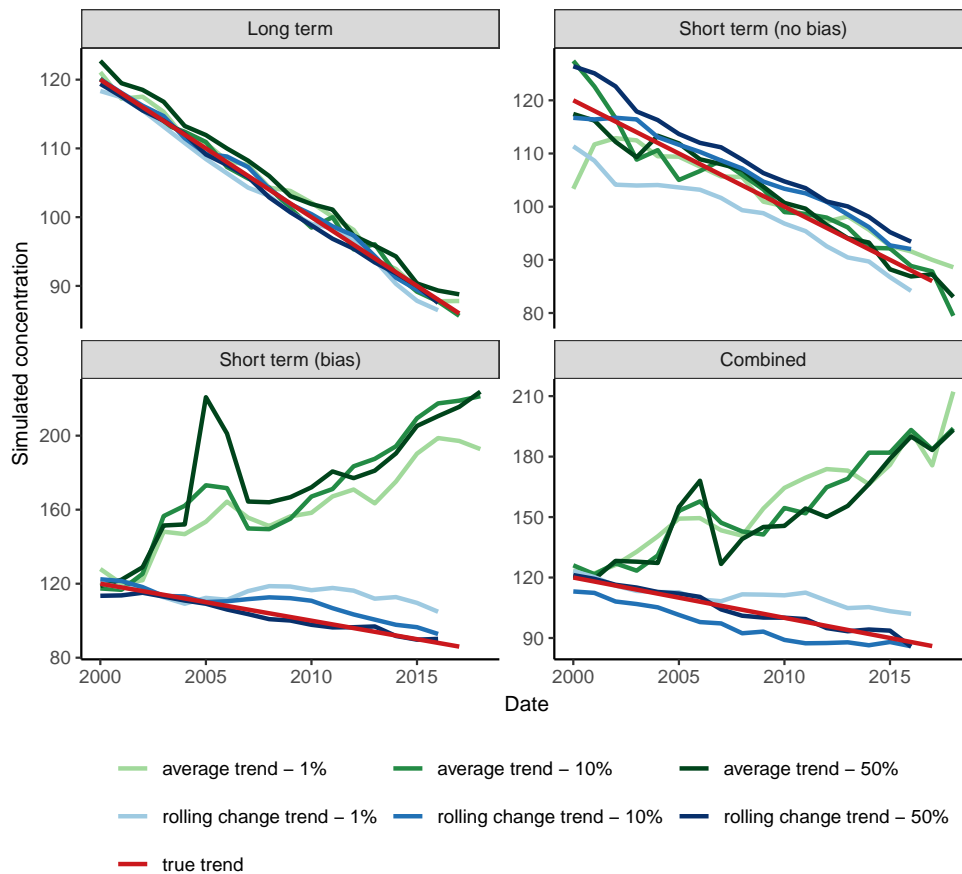


Figure 2.3: Comparison of the average trend and rolling change trend ($n = 3$) with the true trend of simulated data for four different scenarios. In each case, the trends are derived from 100 random samples, each of 100 simulated time series. The lines correspond to the trends with NCC equal to the 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.

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 Figure 2.4. 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.

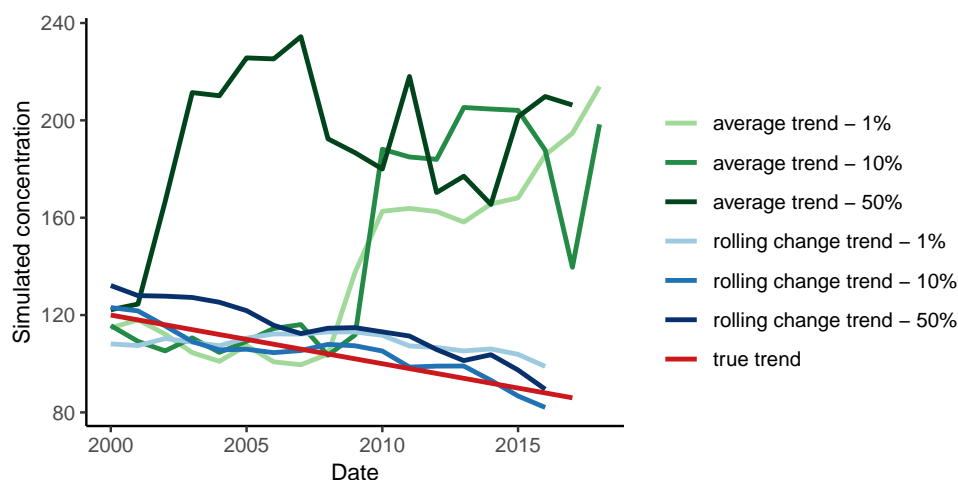


Figure 2.4: Comparison of the average trend and rolling change trend ($n = 3$) with the true trend of data from 4 time series simulated using the ‘combined’ scenario. The trends are derived from 100 random samples of simulated data. The lines correspond to the trends with NCC equal to the 50th, 10th 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.

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 Appendix II (Figure B.1). 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.

2.3.2 Long term trends in London ambient air quality

2.3.2.1 Identification of the bias effect on the trend

Comparison of the average trend over all London roadside sites 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 Figure 2.5. 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 Appendix II, Figure B.2). 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 NO_2 .

The effect of the bias in site location on the trend in average roadside NO_x and NO_2 concentrations can be observed through a comparison of the rolling trends over rolling windows of different widths (n), as shown in Appendix II (Figures B.3 - B.5).

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 Appendix II, Figures B.6 - B.8), in corroboration of the findings of the Defra report discussed in Section 2.1.1 (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 NO_2 concentrations are dominated by non-local background sources rather than local traffic sources, which constitute the major source at roadside sites.

2.3.2.2 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 2.2.2 was applied to the London roadside monitoring data.

The rolling change trends in NO_x concentration, NO_2 concentration and NO_2/NO_x ratio are shown in Figure 2.5 (right). In all cases, these derived trends bear a far closer similarity with the trend for the long-term site (Figure 2.5 (middle)) than with the biased average trend (Figure 2.5 (left)), offering further evidence in support of the technique's efficacy.

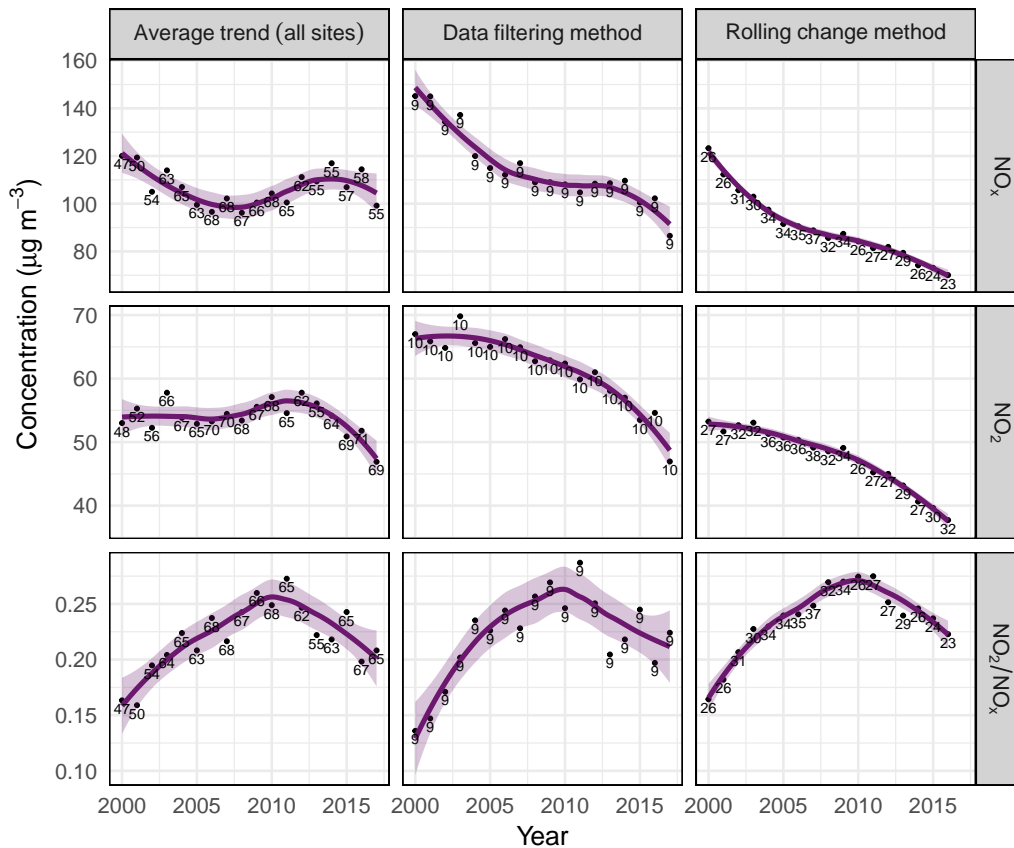


Figure 2.5: Comparison of the rolling change trends in NO_x concentration, NO_2 concentration, and NO_2/NO_x ratio at London roadside sites 2000–2017, using $n = 3$ ('Rolling change method') with the trend in the average concentration using data from (i) all available monitoring sites ('Average trend (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 fit. The numbers at each data point correspond to the number of monitoring sites contributing to the data point.

The rolling change technique reveals a more optimistic trend from 2000–2017 in NO_x concentration at London roadside sites than that implied by the average trend. Table 2.1 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 (iii) the trend derived using the rolling change method).

Application of the Theil-Sen estimator to the NO_x concentration trends in Figure 2.5 yielded a slope of -2.5 [$-3.3, -2.0$] $\mu\text{g m}^{-3} \text{ year}^{-1}$ for the rolling change trend. In contrast, the gradient of the average trend was -0.2 [$-1.2, 1.0$] $\mu\text{g m}^{-3} \text{ year}^{-1}$. 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.

Table 2.1: 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 , NO_2 and NO_2/NO_x concentration at roadside in London 2000-2017.

Pollutant	Method	Theil-Sen slope ($\mu\text{g m}^{-3} \text{ year}^{-1}$)	95% confidence interval
NO_x	Average trend (all sites)	-0.2	[-1.2, 1.0]
NO_x	Average trend (longterm sites)	-2.6	[-3.3, -1.4]
NO_x	Rolling change method	-2.5	[-3.3, -2.0]
NO_2	Average trend (all sites)	-0.1	[-0.5, 0.2]
NO_2	Average trend (longterm sites)	-1.0	[-1.2, -0.6]
NO_2	Rolling change method	-0.9	[-1.1, -0.7]
NO_2/NO_x	Average trend (all sites)	0.0	[0.0, 0.0]
NO_2/NO_x	Average trend (longterm sites)	0.0	[0.0, 0.0]
NO_2/NO_x	Rolling change method	0.0	[0.0, 0.0]

The differences between the average and rolling change trends in NO_2 concentration were less extreme, but nonetheless notable. Theil-Sen slope of the rolling change trend was -0.9 $[-1.1, -0.7] \mu\text{g m}^{-3} \text{ year}^{-1}$ in comparison to -0.1 $[-0.5, 0.2] \mu\text{g m}^{-3} \text{ year}^{-1}$ 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 2008-2012.

The effectiveness with which the rolling change trend represents the ‘true trend’ was evaluated by comparison with trends in NO_x and NO_2 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, 2020) is shown in Figure 2.6. The emissions data shows a monotonic, almost linear downward trend between 2000-2016, similar to the rolling change trend in NO_x concentration from the London data (see Figure 2.5). 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




Figure 2.6: Trend in UK NO_x emissions from road transport (urban driving) sectors between 2000-2016 (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 fit.

sources also contribute. One such source is natural gas combustion for domestic heating, from which NO_x emissions have decreased less between 2000-2017 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, 2020; Wakeling et al., 2018).

A recent study of London air quality using satellite data estimated a trend in NO_2 concentration of $-0.23 \times 10^5 \text{ molecules cm}^{-2} \text{ year}^{-1}$ between 2005-2015, which is approximately $-1.76 \text{ } \mu\text{g m}^{-3} \text{ year}^{-1}$, 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 NO_2 concentration over the same period from roadside monitoring sites was $-1.0 [-1.5, -0.7] \text{ } \mu\text{g m}^{-3} \text{ year}^{-1}$, compared to the average trend slope of $-0.02 [-0.4, 0.3] \text{ } \mu\text{g m}^{-3} \text{ year}^{-1}$. 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 NO_2 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 NO_2/NO_x vehicle *emission ratio* across Europe follows a pattern of increase from 1995-2008 then decrease between 2009-2015. This pattern is replicated in the

NO₂/NO_x rolling change trend shown in Figure 2.5 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 in London between 2005-2009 and 2010-2014 with those obtained from the rolling change trend and the average trend are shown in Figure 2.7. As mentioned in Section 2.1.1, 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.

The roadside increments were calculated using the same background monitoring site as was used in the study by Font and Fuller (2016). There, the North Kensington site was chosen as it had a long time series that was representative of the trends seen in the time series' of other London background sites.

As can be seen in Figure 2.7, 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 2005-2009 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 2010-2014, in corroboration of the findings of Font and Fuller (2016).

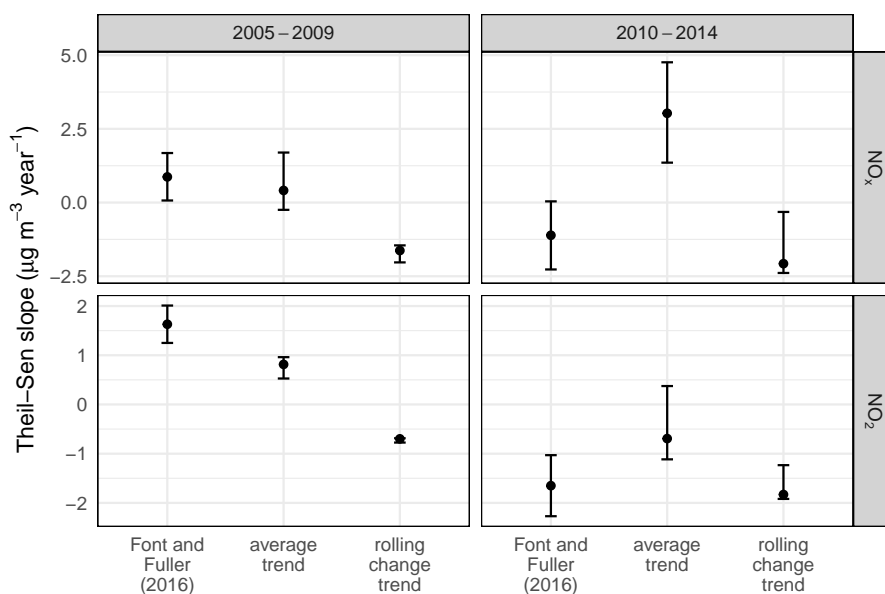


Figure 2.7: Comparison of the Theil-Sen slope calculated by [Font and Fuller \(2016\)](#) with the rolling change trend and the average trend in NO_x and NO₂ roadside increments at London roadside monitoring sites between 2005-2009 and 2010-2014. The error bars represent 95% confidence intervals.

[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 2.3, 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.

2.3.3 Long term trends in ambient air quality in Scotland and the UK

The number of monitoring sites measuring NO_x , NO_2 , and PM_{10} concentration in the Scottish network in each year between 2000 and 2017 is shown in Figure 2.8. It can be seen that the number of sites increases over time for all pollutants and, as a result, the trend in average concentration over all available data would be an inappropriate method to visualise the long term trends in ambient concentration for the reasons outlined in Section 2.1.1.




Figure 2.8: The number of roadside monitoring sites measuring NO_x , NO_2 and PM_{10} concentration in Scotland by year between 2000 and 2017.

Furthermore, Figure 2.9 demonstrates the problems inherent in the application of data filtering methods to derive a long term trend that is unaffected by site flux within the monitoring network. The map on the left shows all roadside monitoring sites measuring NO_2 concentration at any time between 2000 and 2017, while the map on the right shows the sites measuring NO_2 concentration constantly over the duration of the period ('long term sites'). By using data filtering to resolve problem of site flux, the data that can be included in the trend analysis is limited to that measured at only one monitoring site, and therefore subject to the same problems discussed in Section 2.1.1 when conducting trend analysis on data from

a single site: namely, that the trend is unlikely to be representative of air quality at a regional or national scale, as well as strongly influenced by local variations.

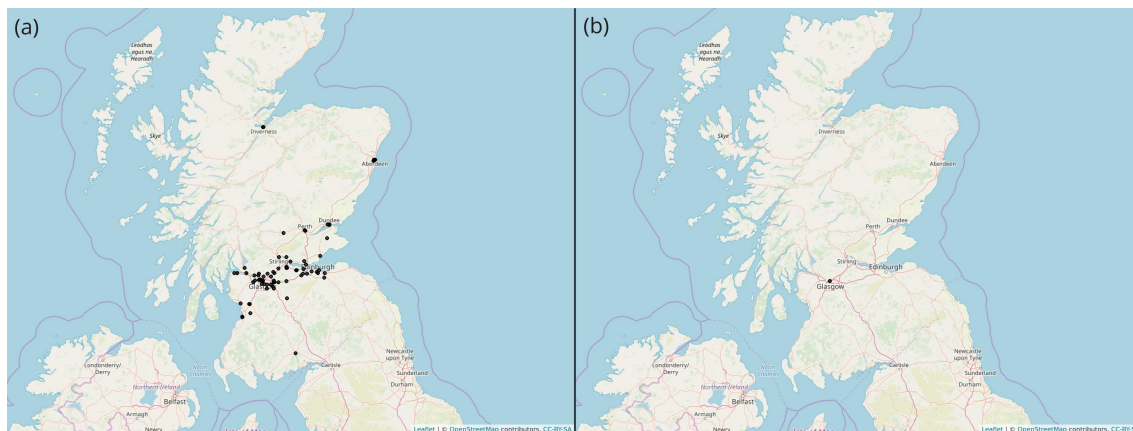


Figure 2.9: The spatial distribution of roadside monitoring sites measuring roadside NO_2 concentration in Scotland (a) at any point between between 2000 and 2017 (left), and (b) constantly throughout the duration of the period 2000-2017 (right).

Figure 2.10 shows the trends in roadside NO_2 concentration in Scotland between 2000 and 2017, calculated using the three methods mentioned previously. The top-left plot shows the trend in the average concentration using data from all available monitoring sites, the top-right plot shows the trend using data from the single long term monitoring site (as shown in Figure 2.9), and the bottom plot shows the trend calculated using the rolling change method. In each plot, the numbers beside each data point represent the number of monitoring sites contributing data to the average.

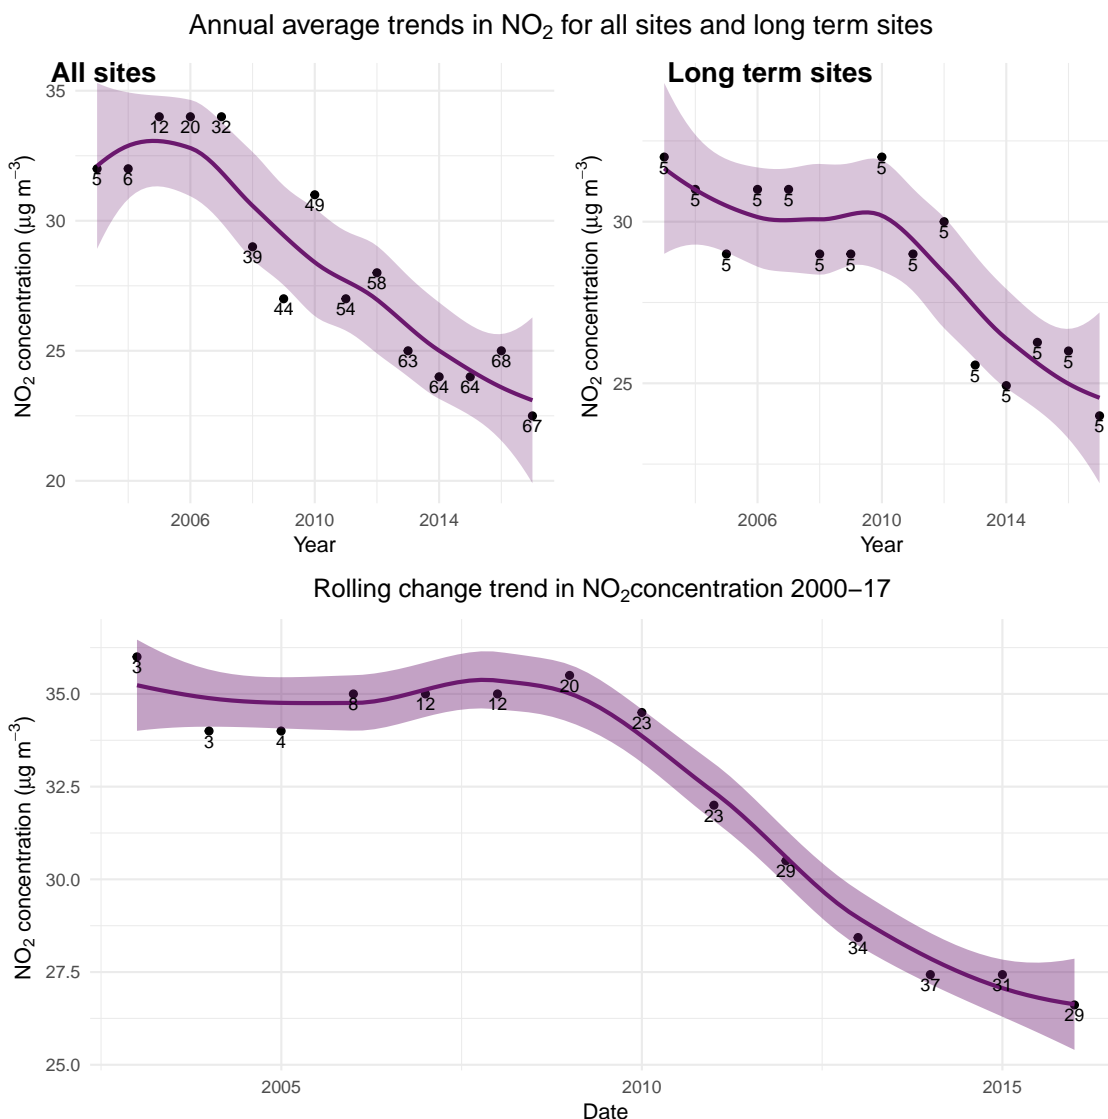


Figure 2.10: The trends in roadside concentration of NO₂ in Scotland between 2000 and 2017 calculated using three different methods: (i) the average concentration across all available data (top left), (ii) the average concentration using data only from sites measuring constantly throughout the analysis period (top right), and (iii) the rolling change method (bottom). The smoothed lines are loess (local regression) fits, with the 95% confidence interval represented by the shaded band. The numbers signify the number of monitoring sites contributing to each annual data point.

It is clear that each method results in a very different trend, and therefore the conclusions drawn regarding the changes in air quality over this period would differ considerably depending on the choice of method for calculating the trend. While the trend in the average concentration suggests that after an initial decline in NO₂ concentrations until 2004, the rate of decrease in concentration slowed and has changed little since 2004, the rolling change method trend reveals that, in actuality, roadside NO₂ concentrations changed little between 2000 and 2010, and

since then have declined monotonically. A more comprehensive exploration of the ability of the rolling change method to more accurately represent the trend is provided in Lang et al. (2019). Figure 2.10 also demonstrates that the rolling change method enables the retention of more data in the trend analysis than is possible using the data filtering method, as can be seen from the numbers beside the data points.

The rolling change method was applied to data from the Scottish monitoring network to calculate the national-scale trends in roadside NO_x , NO_2 , and PM_{10} concentrations, as well as the NO_2/NO_x ratio, in Scotland between 2000 and 2017. These trends are shown in Figure 2.11.

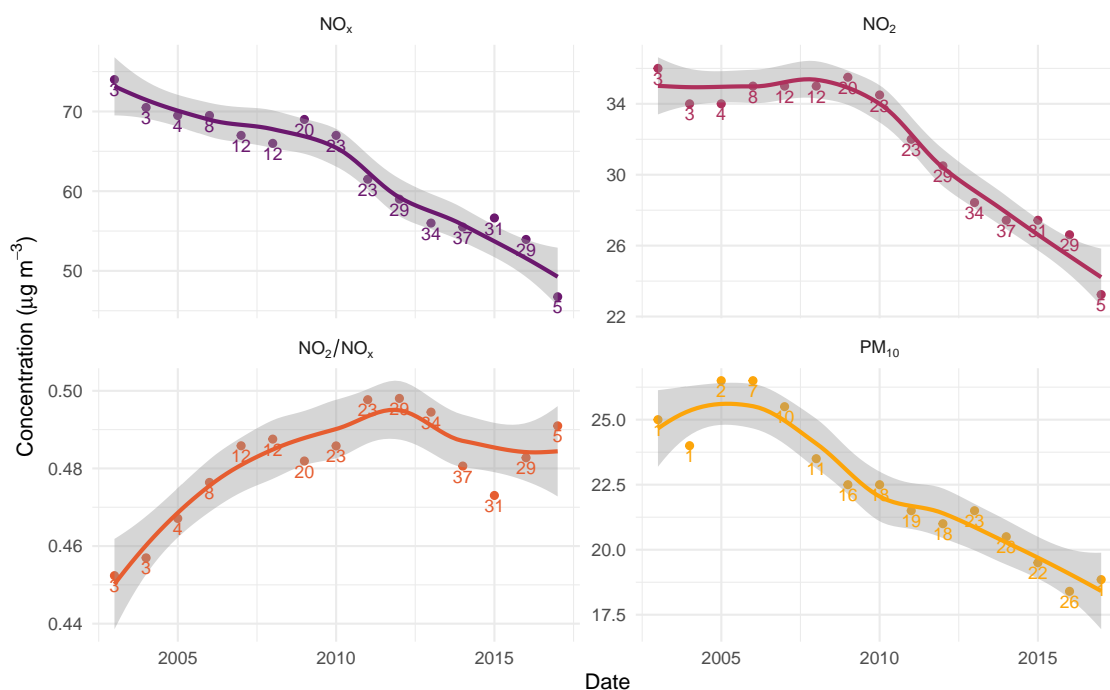


Figure 2.11: Annual rolling change trends in NO_x , NO_2 , NO_2/NO_x and PM_{10} concentrations in Scotland 2003-2018. The numbers indicate the number of monitoring sites contributing to each (annual) data point.

The roadside NO_x concentration is observed to have decreased monotonically since 2002. As mentioned previously, the NO_2 concentration was initially stable, until it began to decrease monotonically in 2010. These changes can be attributed to the introduction of vehicle exhaust technologies aimed at reducing emissions of these species, such as three-way catalysts used on petrol vehicles and the more recent use of Lean NO_x Traps (LNT) and Selective Catalytic Reduction (SCR) on diesel vehicles over this period.

The NO_2/NO_x ratio was observed to increase between 2003 and 2011, before

reaching a turning point, and decreasing thereafter. The initial increase in the NO_2/NO_x ratio stems from the introduction of vehicle emission after-treatment technologies, such as Diesel Oxidation Catalysts (DOC) and Diesel Particulate Filters (DPF), which deliberately oxidise NO to NO_2 for use in the oxidation of other pollutants, such as CO, hydrocarbons and particulate matter. While Figure 2.11 demonstrates that the ambient concentration of NO_x has decreased since 2002, the introduction of these vehicle emission technologies resulted in an increase in the proportion of NO_x emitted as NO_2 from vehicle exhaust, resulting in an increase in the NO_2/NO_x ratio.

The reasons for the observed decline in NO_2/NO_x ratio since 2011 are less clear, and several factors have likely contributed. Vehicle emission remote sensing measurements have found that the NO_2/NO_x ratio decreases for diesel passenger vehicles as the mileage increases [Carslaw et al. \(2019\)](#). Additionally, it is likely that vehicle emission after-treatment systems have been modified to no longer over-produce NO_2 .

The PM_{10} concentration was relatively constant between 2000 and 2006, before declining monotonically for the rest of the period analysed. This decline can be linked to the fitting of many Euro 4 vehicles, and all post-Euro 5 vehicles, with diesel particulate filters (DPF).

The trends in air quality in Scotland (Figure 2.11) were compared to the trends calculated using the same methodology for the UK and Europe over the same period. These trends are shown in Figures 2.12 and 2.13 respectively.

CHAPTER 2. A NEW TREND ANALYSIS APPROACH FOR AIR QUALITY NETWORK DATA

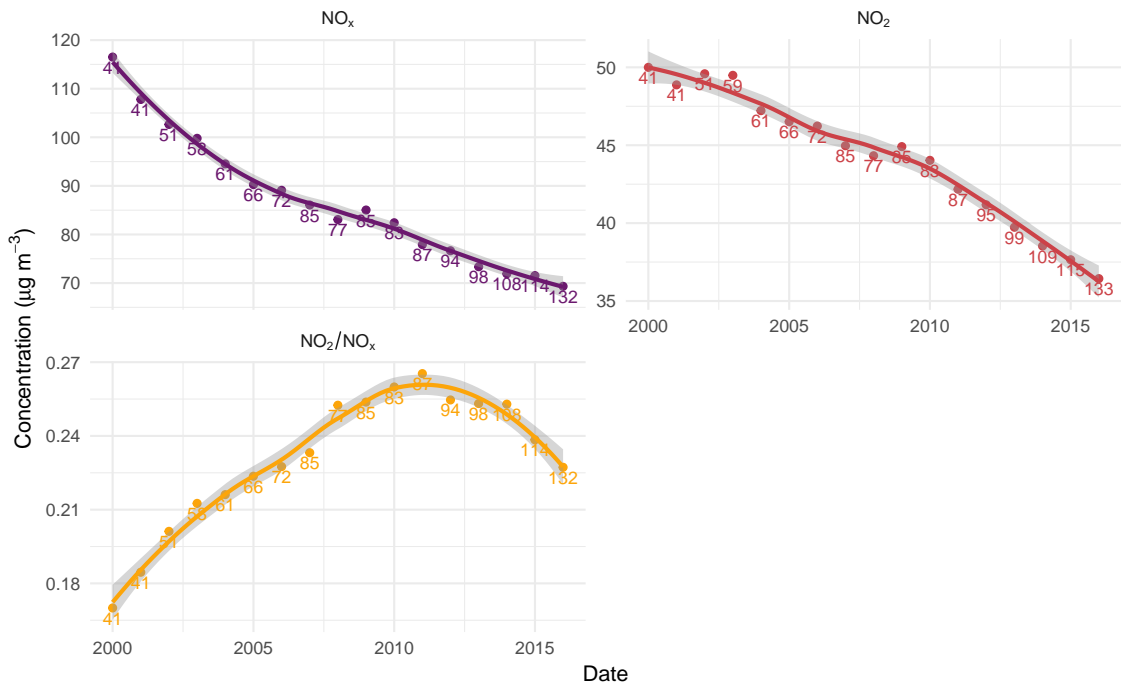


Figure 2.12: Annual rolling change trends in NO_x and NO_2 concentration and NO_2/NO_x ratio at roadside in the UK 2000-2017. The numbers indicate the number of monitoring sites contributing to each (annual) data point.

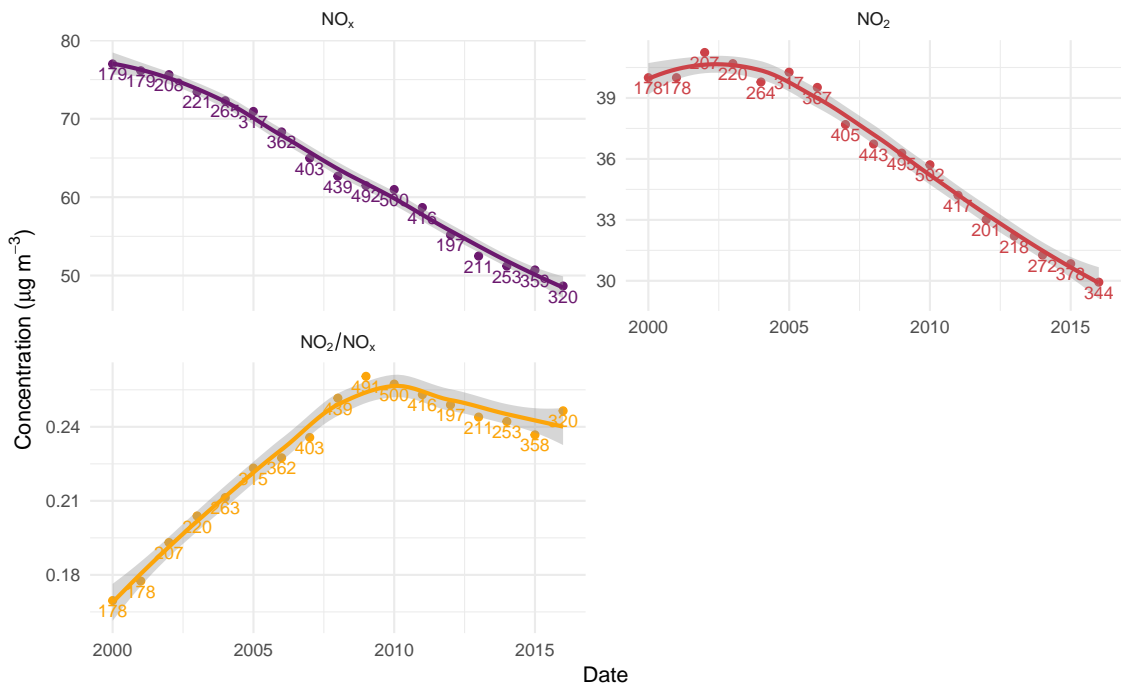


Figure 2.13: Annual rolling change trends in NO_x and NO_2 concentration and NO_2/NO_x ratio at roadside in Europe 2000-2017. The numbers indicate the number of monitoring sites contributing to each (annual) data point.

The trends in Figures 2.12 and 2.13 share similar shapes to the trends in Scotland, corroborating the supposition that these aggregate trends represent the large-scale (i.e. national) changes in air quality. The drivers of these national-scale changes, for example, changes in vehicle emission technologies, are likely to be common to both Scotland and Europe, while the aggregation of the data on a national scale eliminates the influence of more local factors, such as local policy changes and urban/environmental changes, that may otherwise obscure these effects. It is noticeable that the turning points in the trends in NO₂ and PM₁₀ concentrations and the NO₂/NO_x ratio in Scotland lag several years behind their counterparts in Europe.

The observed trends in NO_x and NO₂ concentrations, and NO₂/NO_x ratio are consistent with the findings of other studies. A study by Grange et al. (2017) found that, in Europe, NO_x concentrations decreased between 1998 and 2015, and that NO₂ concentrations increased between 1997 and 2009, before decreasing until 2015. Grange et al. (2017) also showed that the NO₂/NO_x ratio in Europe had increased between 1995 and 2008, and then decreased between 2009 to 2015. These results corroborate the trends calculated using the rolling change method, although it is not clear why the trends in Scotland appear to lag behind those in Europe by a couple of years. Plausible explanations include differences in the composition of the vehicle fleet (for example, a greater proportion of diesel vehicles), or an older vehicle fleet, however we lack the high resolution vehicle fleet data for Scotland and Europe to investigate further.

2.3.4 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 2.1.1, 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 2000-2017 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 NO_2 concentration over the period 2000-2017, none of which has more than 1 monitoring site. 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 usable 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.

2.4 Conclusions

Long term trend analysis is an important tool for measuring changes in air quality over time, and evaluating the effects of policy interventions on ambient air quality. In order to evaluate the changes in air quality on a large-scale (for example, a regional, national, or continental scale), it is necessary to aggregate data from multiple monitoring sites to 'average out' the effects of local variability.

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 with different durations. Site flux, that is, 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 in order 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, which we call the 'rolling change method', 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 rolling change method was applied to an analysis of the long term trends in air quality at the roadside in London, Scotland, the UK and Europe. These trends provided a generally positive view of changes in air quality. In each case, concentrations of NO_x have decreased monotonically. In London and the UK, NO_2 concentrations have also decreased monotonically over the entire period, while in Scotland and Europe were stable until 2010, in the case of the former, and 2003 in the case of the latter, before declining. In each case, the NO_2/NO_x ratio increased to a peak before declining. These changes were explained in terms of changes in vehicle emission technologies that have occurred over the period, notably, the introduction of DOC and DPF. PM_{10} concentrations in Scotland, the UK and Europe were also shown to have decreased monotonically since 2006, and this was attributed to the introduction of DPF to most Euro 4 vehicles, and to all Euro 5 vehicles. The similarity in the shape of the trends on different scales indicated that they shared common drivers, and therefore were attributable to large-scale changes rather than local variation.

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.

Bibliography

- Anttila, P., Tuovinen, J.-P., 2010. Trends of primary and secondary pollutant concentrations in Finland in 1994–2007. *Atmospheric Environment* 44 (1), 30–41.
URL <https://doi.org/10.1016/j.atmosenv.2009.09.041>
- Carslaw, D. C., Farren, N. J., Vaughan, A. R., Drysdale, W. S., Young, S., Lee, J. D., 2019. The diminishing importance of nitrogen dioxide emissions from road vehicle exhaust. *Atmospheric Environment: X* 1, 100002.
URL <https://doi.org/10.1016/j.aeaoa.2018.100002>
- Defra, 2014. Air Pollution in the UK 2013. Tech. rep., Department for Environment, Food and Rural Affairs.
URL <https://uk-air.defra.gov.uk/library/annualreport/viewonline?year=2013{ }issue{ }1{#}report{ }pdf>
- Duyzer, J., van den Hout, D., Zandveld, P., van Ratingen, S., 2015. Representativeness of air quality monitoring networks. *Atmospheric Environment* 104, 88–101.
URL <https://doi.org/10.1016/j.atmosenv.2014.12.067>
- Fleming, Z. L., Doherty, R. M., Von Schneidemesser, 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., 2018. Tropospheric Ozone Assessment Report: Present-day ozone distribution and trends relevant to human health. *Elementa: Science of the Anthropocene* 6 (1), 12.
URL <http://www.elementascience.org/article/10.1525/elementa.273/>
- Font, A., Fuller, G. W., 2016. Did policies to abate atmospheric emissions from traffic have a positive effect in London? *Environmental Pollution* 218, 463–474.
URL <https://doi.org/10.1016/j.envpol.2016.07.026>
- Grange, S. K., Lewis, A. C., Moller, S. J., Carslaw, D. C., 2017. Lower vehicular primary emissions of NO₂ in Europe than assumed in policy projections. *Nature*

BIBLIOGRAPHY

- Geoscience 10 (12), 914–918.
URL <http://www.nature.com/articles/s41561-017-0009-0>
- Guerreiro, C. B., Foltescu, V., de Leeuw, F., 2014. Air quality status and trends in Europe. *Atmospheric Environment* 98, 376–384.
URL <https://doi.org/10.1016/j.atmosenv.2014.09.017>
- Lang, P., 2018. aqtrends.
URL <https://github.com/pollylang/aqtrends>
- Lang, P. E., Carslaw, D. C., Moller, S. J., 2019. A trend analysis approach for air quality network data. *Atmospheric Environment: X* 2, 100030.
URL <https://doi.org/10.1016/j.aeaoa.2019.100030>
- Malley, C. S., Von Schneidemesser, 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. *Atmospheric Chemistry and Physics* 18, 3563–3587.
URL <https://doi.org/10.5194/acp-18-3563-2018>
- Masiol, M., Squizzato, S., Formenton, G., Harrison, R. M., Agostinelli, C., 2017. Air quality across a European hotspot: Spatial gradients, seasonality, diurnal cycles and trends in the Veneto region, NE Italy. *Science of The Total Environment* 576, 210–224.
URL <https://doi.org/10.1016/j.scitotenv.2016.10.042>
- Mavroidis, I., Chaloulakou, A., 2011. Long-term trends of primary and secondary NO₂ production in the Athens area. Variation of the NO₂/NO_x ratio. *Atmospheric Environment* 45 (38), 6872–6879.
URL <https://doi.org/10.1016/j.atmosenv.2010.11.006>
- NAEI, 2020. UK National Atmospheric Emissions Inventory.
URL <http://naei.beis.gov.uk/index>
- Pope, R. J., Arnold, S. R., Chipperfield, M. P., Latter, B. G., Siddans, R., Kerridge, B. J., 2018. Widespread changes in UK air quality observed from space. *Atmospheric Science Letters* 19 (5), 1–8.
URL <http://doi.wiley.com/10.1002/asl.817>
- 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.,

BIBLIOGRAPHY

Watterson, J., Pearson, B., Rushton, K., 2018. UK Informative Inventory Report (1990 to 2016). Tech. rep., NAEI.

URL <http://ec.europa.eu/environment/air/pollutants/ceilings.htm>

Wang, J., Zivot, E., 2006. Rolling Analysis of Time Series. In: Modeling Financial Time Series with S-PLUS®. Springer New York, New York, NY, Ch. 9, pp. 313–360.

URL http://link.springer.com/10.1007/978-0-387-32348-0_{_}9

3. Development and Application of Random Forest Models to Air Pollutant Time Series

3.1 Introduction

3.1.1 Background

In response to frequent exceedances of WHO and EU limits on ambient concentrations of major air pollutants, particularly NO₂ and PM, many European countries have implemented policy interventions aimed at reducing air pollution in urban areas. These interventions typically target the road traffic source, and include low emission zones (LEZ), clean air zones (CAZ), variable speed limits, and reduction in speed limits, as well as other mechanisms for reducing congestion and stop-start traffic flows.

Such policies are not without their costs, both to the road user in terms of the costs in time, money and convenience, and to the government. It is therefore important to estimate the effectiveness of the intervention in achieving its goals in order to justify its existence to the public. During the planning stage of the intervention, the benefits of the measure are estimated using models, however these estimates only predict the potential benefits of a perfectly implemented system. To evaluate the actual effects of the intervention on ambient air quality, it is necessary to analyse real-world monitoring data.

3.1.2 Intervention and Accountability Studies

Intervention analysis was previously discussed in Chapter 1. It concerns the investigation of the effect of policy interventions on air pollutant emissions, air quality, exposure and, ultimately, human health. The main challenges faced by accountability studies are the establishment of an appropriate control scenario, and the problem of distinguishing the changes in air pollutant concentration resulting from the policy intervention from the changes due to confounding factors such as meteorology.

This chapter presents a method that addresses both of these issues. Random forest models, with an indicator variable that represents the phase of the LEZ, are used to remove the effects of confounding factors, and to generate a control scenario for comparison.

3.1.3 Random Forest Modelling

Random forest is an ensemble decision tree model, which has gained popularity recently due to its accuracy in a range of domains combined with its excellent

interpretability (Donges, 2019; Sarica et al., 2017).

Decision trees are a powerful method for modelling non-linear relationships. However, a single decision tree suffers from a common dichotomy in machine learning: the bias-variance trade-off. A decision tree is a non-parametric technique, meaning that it learns the distribution of the data directly from the data, rather than relying on over-simplistic a priori assumptions, such as linearity. Non-parametric methods provide increased flexibility to the model, often resulting in improved accuracy, however it also introduces the risk of 'overfitting'. Overfitting is a common problem in machine learning applications, and refers to the situation where an over-complex model is fit to the noise in the training data, rather than the signal. The trained model, therefore, describes the training data very well, but performs poorly when generalising to new cases. If the decision tree is grown too deep, it risks overfitting to the data, even to the extreme case where each node of the tree contains only a single data point (high variance). However, too shallow a tree sacrifices accuracy, particularly in data sets containing many predictors (high bias). Ensemble methods, such as random forest, ameliorate the effects of overfitting in order to improve the accuracy of the model predictions by aggregating the predictions of many small models.

The ensemble method used in random forest is bagging, which is short for 'bootstrap aggregation'. Bagging improves the accuracy of the ensemble by bootstrapping (generating a new data set by drawing samples from the original data set with replacement) the training data, and using each sample to train a shallow decision tree. Additionally, for each tree the features themselves are also bootstrapped - typically a sample of $m = p/3$ for regression trees (where p is the total number of features) is used to train each tree. This has the effect of de-correlating individual trees, which reduces the variance of the individual model and makes the ensemble robust to correlated predictors. Once all of these small, independent trees are trained, they predict the response of new input data by voting, i.e. the average of the predictions of all the trees is the overall outcome.

By using ensemble methods to combine the advantages of decision trees with a robustness to overfitting, the random forest inherits an unmatched ability to model non-linear relationships, and a robustness to correlations between predictor variables from the decision tree method. This makes it very useful for problems involving non-linear relationships and interaction effects, such as are common in air quality data. Additionally, random forest is robust to the presence of uninformative features, and handles missing values and outliers well.

In the context of the application to air pollutant concentrations, these methods

help to capture important characteristics. For example, it is well-established that the relationship between a variable such as wind speed and pollutant concentration is non-linear. Furthermore, different meteorological variables do not act independently of one another i.e. there is an *interaction* between them. For example, the relationship between a pollutant concentration and wind speed is not the same for all wind directions or ambient temperatures. Tree-based methods can account for both non-linearity and interactions in a way that requires no prior knowledge of their functional relationships. Moreover, for changes that occur abruptly, tree-based approaches are able to deal with these situations, which might prove more challenging for other techniques that use smoothing, such as Generalized Additive Models. In this respect, they are potentially well-suited to intervention analysis where changes might occur over short time periods.

3.1.4 Model Interpretation

The uptake of complex machine learning models, such as artificial neural networks, has been impeded in large part due to their lack of interpretability, leading them to be labelled 'black boxes'. An attractive feature of random forests is that it is more interpretable than most machine learning models. The model produces estimates of the relative importance of the predictor variables, partial dependence plots and interaction plots. The interpretability of the model is important because it helps to establish whether the relationships between variables are chemically and physically plausible.

3.1.4.1 Variable Importance

The variable importance is calculated using two methods (in most implementations):

1. Accuracy-based — the decrease in the accuracy of the model when the variable is removed is estimated. For each tree, there is an out-of-bag subset of data that the model was not trained on that is used for evaluating the model performance. First, the predictive accuracy of the model is calculated using the out-of-bag data set, then the values of the variable are randomly shuffled and the accuracy is re-calculated. The randomly shuffled variable is assumed to have no predictive power, so the shuffling is equivalent to removing the variable. The variable importance is estimated as the mean decrease in predictive accuracy across all trees.

2. Gini-based — The Gini index is used in node splitting as a measure of node purity. It can also be used to measure variable importance, by calculating the sum of the Gini decrease (how much the Gini index decreased from the parent node to the sub-node) at all nodes where the variable was used to split the data, across all trees in the forest. The sum is divided by the number of trees in the forest to give an average (Hoare, 2019; Lee, 2017).

3.1.4.2 Partial Dependence Plots

Partial dependence plots can be considered as analogous to variable coefficients in multiple linear regression. They are functions that describe the relationship between each variable and the response (dependent variable) if all other variables in the model are kept the same. However, unlike linear regression variable coefficients, partial dependence plots can describe a non-linear relationship between variables. Partial dependence plots are extremely useful in practise, because it gives the user insight into how the model is making its predictions, as well as ascertaining the relationship between the response variable and a predictor (for example, the intervention indicator) in the absence of effects from other predictors in the model (e.g. the confounders).

The partial dependence function is computed using the following algorithm:

1. Find the unique values of the variable of interest in the training data set.
2. Create one replicate of the training data set for each unique value of the variable of interest. Fix the value of the variable of interest to a single unique value for each replicate.
3. For each replicate, predict the value of the response for each observation using the model.
4. For each replicate, average the predicted values of the response. This average is the value of the response mapped from the unique value of the variable of interest for this replicate in the partial dependence function (Wright, 2018).

Partial dependence plots are particularly useful for understanding whether the dependent variable varies in ways that can be understood by the underlying processes of dispersion and atmospheric chemistry. For example, for a primary pollutant ground-level source, it would generally be expected the concentration would decrease with increasing wind speed. Similarly, as ambient temperature increases, it would be expected that concentrations of pollutants would decrease

owing to the enhanced thermal turbulence. Clearly, the actual variations would depend on the situation in question, but the ability to critically evaluate the responses in this way is a valuable characteristic of tree-based models.

3.1.5 Meteorological Normalisation

Meteorological normalisation is a process whereby the effects of meteorology on the ambient concentration are removed in order to more clearly quantify changes resulting from other factors such as changes in emission source strength. This is done by calculating a ‘normalised trend’: the trend under ‘average’ conditions of all predictor variables other than the date (which represents the long term trend). The ultimate result is that the effects of variation in the model predictors (e.g. meteorology and seasonality) are removed from the time series, leaving only the long term variation (i.e. the long term trend) (Grange et al., 2018).

The calculation of the normalised trend is similar to that of the more familiar partial dependence plots. A previously trained random forest model is used to predict each pollutant concentration (i.e. each value of the response) 500 times. Each time, the values of the predictors other than date are randomly sampled with replacement from the data set. The arithmetic mean of the 500 predictions is then calculated, and assumed to represent the concentration of the pollutant under ‘average’ conditions of the modelled variables. The full details of the method are described by Grange et al. (2018) in their original paper.

3.1.6 Chapter Summary

The analysis of the use of random forest models and the meteorological normalisation methodology to conduct intervention analyses followed three main themes:

- Method validation — testing the efficacy of the method by modelling pollutants with well-known long term trends.
- Method analysis — investigation of variations of the method to achieve optimal performance
- Method application — application of the method to real-world air quality problems (primarily intervention analyses)

3.2 Methods

All analysis was carried out in R. The code used for the random forest modelling can be found in the *rfmodels* package (Lang, 2018). The code used to carry out meteorological normalisation is based on the normalisation functions in the *rmweather* package (Grange, 2018).

3.2.1 Data Preparation

Routine ambient monitoring data were collected from the Automatic Urban and Rural Network (AURN) maintained by Defra, the London Air Quality Network (LAQN) run by King's College London, the Air Quality England (AQE) database collected by Ricardo Energy & Environment, the Scottish Air Quality Network (SAQN), and the Welsh Air Quality Network (WAQN) (Department for Environment, Food and Rural Affairs (Defra), 2019; Kings College London, 2019; Ricardo Energy & Environment, 2019). The AQE, SAQN and WAQN databases store data collected by local authorities in England, Scotland and Wales, respectively. The raw data consisted of hourly observations of air pollutant concentrations.

The hourly concentration of the pollutant of interest measured at the monitoring site under analysis were used as the response variable in the model. The explanatory predictor variables included background concentrations of the pollutant of interest, surface meteorology, traffic counts and temporal variables.

The background concentrations were the data measured at the nearest urban background site to the monitoring site of interest which measured the pollutant of interest with sufficient data capture (in this case, 80%) over the period of interest.

Surface meteorological measurements (air temperature, wind direction, wind speed, atmospheric pressure, relative humidity, dew point and visibility) were sourced from the Integrated Surface Database (ISD) (NOAA, 2016). The monitoring site from which to join the meteorological measurements was chosen in the same way as the background site: the closest meteorological station with sufficient data capture for the selected meteorological variables was chosen.

Traffic count data was sourced from the Department for Transport's annual road traffic counts (Department for Transport, 2010). Data was collected from the road link on which the monitoring site was located. The annual count data was joined to the hourly monitoring and meteorological data by assigning the annual count to every observation within that year.

Missing values in all variables were interpolated with the median concentration.

3.2.2 Model Optimisation, Training and Evaluation

Random forest requires several parameters to be set by the user prior to model training. The optimal values of the parameters were ascertained by grid search over a range of possible values. The parameter combination with the lowest OOB RMSE (out-of-bag root mean squared error) was selected. The tuned parameters were:

- The number of variables to randomly sample at each split (i.e. the number of variables each tree is trained on) was set to 7.
- The number of observations to train on was set to 85%. This is key to the bias-variance trade-off of the model: lower values reduce training time but introduce more bias, and lower values increase variance and risk overfitting to the training data.
- The minimum number of observations in the terminal nodes (the node size) was set to 3. This parameter represents the minimum depth of the trees, and controls the complexity of the trees. Smaller values lead to deeper, more complex trees, while larger values lead to shallow trees. It also affects bias-variance trade-off: shallow trees introduce more bias, while deeper trees introduce more variance and risk overfitting.

The random forest model was trained on the hourly ambient air pollutant concentration (at the roadside monitoring site) as the response, and the predictors: meteorological data (e.g. air temperature, wind speed, wind direction, relative humidity, atmospheric pressure, dew point, visibility), background concentrations (hourly air pollutant concentration measurements from a nearby representative urban background monitoring site), traffic flow data (annual traffic counts), temporal variables (hour of day, day of week, Julian day) and the long term trend (Unix date). Additional predictor variables may also be included as required. For example, in the intervention analysis in Section 3.3.4, an intervention indicator variable is often included to model the presence/absence of the intervention.

In the context of an intervention analysis, a temporal control can be established by training the model using data from before and after the implementation of the intervention. An indicator variable is included in the model, taking the value 0 for data measured when the intervention is not in effect, and 1 when it is in effect.

Three metrics were used to evaluate the performance of the trained models. The R-squared value and the RMSE (Root Mean Squared Error) were calculated using the out-of-bag (OOB) data, and the RMSE was also calculated using the held-out test data set. ‘Held-out data’ refers to data that was not used to train the model, and is therefore unseen by the model. This data is used to evaluate the accuracy of the model in preference to the training data, as it tests the model’s ability to generalise to unseen cases, and potentially enables overfitting to be identified (James et al., 2013).

R-squared is defined as the proportion of the total variation in the response variable that is explained by the model. RMSE is the square root of the variance of the model residuals, and is a measure of the similarity of the predicted values to the observed values of the response (see Equation 3.1) (James et al., 2013; Swalin, 2018).

$$RMSE = \sqrt{\frac{1}{n} \sum_{j=1}^n (y_j - \hat{y}_j)^2} \quad (3.1)$$

3.2.3 Trend Normalisation and Analysis

The trained model was used to calculate the normalised trend in the air pollutant concentration, using the normalisation methodology described in Section 3.1.5.

Trend analysis was conducted on the normalised trend to measure the long-term changes in air pollutant concentration, without the obscuring effects of meteorology, transport and dispersion. Break point analysis was applied to detect change points and, where possible, additional information (e.g. dates of implementation of policy interventions, and the dates from which new vehicle emission standards were applied to new registrations) was considered to inform speculation about their possible drivers (Bai, 1994).

For intervention analyses which included an indicator variable, the effect of the intervention was quantified by calculating partial dependence plots for the indicator variable, or by setting the indicator to different values and comparing the normalised trends.

3.3 Results and Discussion

3.3.1 Method Validation

In order to test whether or not the model is generating an ‘accurate’ normalised trend (i.e. truly representing the trend in concentration with the influence of confounders removed), the analysis was conducted at a location and for a pollutant with a very well understood long term trend: black carbon (BC) at Marylebone Road in London (MY1) between 2000 and 2017. It would be expected that black carbon concentrations at this busy roadside location would be strongly influenced by the vehicles using Marylebone Road, with only minor contributions from other sources, such as wood burning.

The monitoring station at Marylebone Road is situated unusually close to a busy road, and the air quality at the site is therefore overwhelmingly dominated by primary emissions from road traffic. The major drivers of changes in black carbon over this period are understood to be reductions in vehicle exhaust particle emissions as a result of fitting particulate filters. Most post-Euro 4 vehicles, and all post-Euro 5 vehicles are fitted with particulate filters, therefore we would expect to observe reductions in black carbon concentration around 2011 (Euro 5 standards apply to all new vehicle registration).

The normalised trend in black carbon concentration at Marylebone Road in London (MY1) between 2000 and 2018 is shown in Figure 3.1.




Figure 3.1: Monthly average normalised trend in black carbon concentration at Marylebone Road, London 2008-2018. The vertical dashed lines indicate the break points (as calculated using the regression model break point analysis described by Bai (1994)).

Break point analysis indicated that the trend exhibited change points in 2011 and 2015. The change point in 2011, which marks the beginning of a monotonic decline in the black carbon concentration, coincides with the application of the Euro 5/V emission standard to all new vehicle registrations which, crucially, included the fitting of particle filters to all new vehicles. The monotonic decrease between 2011 and 2015 was likely driven by fleet turnover resulting in penetration of Euro 5/V vehicles fitted with particulate filters into the vehicle fleet, replacing more polluting older models. This decrease in black carbon ends at the change point in 2015, and is followed by a period of no change in black carbon concentration. While this change point is coincident with the introduction of the Euro 6/VI emission standard in 2015, the limits on particulate emissions for passenger vehicles and light-duty vehicles did not change between the Euro 5 and Euro 6 standards (since the use of particle filters had reduced these emissions to near-zero) (Williams and Minjares, 2016). Furthermore by 2015 it is probable that most on-road vehicles were Euro 4 standard or higher, and therefore fitted with particle filters, therefore further fleet turnover yields little improvement in air quality.

The normalised trend in black carbon concentration at London Marylebone Road is consistent with expectation, based on the changes in vehicle emission technologies over the period of analysis. This suggests that the method was successful in extracting the long-term trend from the effects of the confounding variables, since if these influences had not been removed, it is likely that the

relatively small changes described here would have been obscured by the much greater variation due to meteorology and dispersion. Further application of the methodology to other air pollutants and locations can therefore be made with greater confidence in the validity of the results.

3.3.2 Method Analysis

The goal of the random forest normalisation method is to produce a normalised trend in the pollutant concentration, where the variation in concentration resulting from variation in meteorology and dispersion (the noise) is minimised, in order to resolve the long term variation in concentration (the signal). One way of achieving this aim is to maximise the accuracy of the model: a more accurate model can account for, and therefore remove, more of the variation due to meteorology from the normalised trend. However, when handling large volumes of data (e.g. hourly monitoring data) and iterative tasks (e.g. sampling during normalisation), the computational demand of the task must also be considered and, if necessary, balanced against model performance.

The random forest methodology comprises a workflow: from data collection, preparation, model training and evaluation, trend normalisation and, finally, trend analysis. Here, two modifications to the workflow are presented. The first aims to maximise the model accuracy by evaluating the relative performance of random forest models trained using meteorological data from different sources. The second aims to minimise the computational demand of the process by estimating the normalised trend using partial dependence plots, rather than carrying out the computationally expensive normalisation procedure.

3.3.2.1 Choice of meteorological data source

In order to evaluate the relative performance of RF models trained using meteorological data from different sources, meteorological data was collected from the following sources:

- Surface measurements of air temperature, wind direction, wind speed, atmospheric pressure, relative humidity, dew point and visibility from the Integrated Surface Database (ISD) ([NOAA, 2016](#))
- Hourly ethane concentration data measured at the London Marylebone Road monitoring site (as a tracer of atmospheric dispersion)

Surface meteorological measurements are the standard data source for modelling ambient air pollutant concentrations (Grange et al., 2018; Dijkema et al., 2008). Its use as a source of meteorological data is self-explanatory.

Ambient ethane concentration was investigated as an alternative source of meteorological information, as it was hypothesised that it would behave as a tracer of atmospheric dispersion. London is built upon a network of natural gas pipelines, which supply the energy needs of the city. This network contains many tiny leaks, which produce a continuous and relatively spatially homogeneous leakage of ethane across the city. As one of the few atmospheric species without a traffic source, the ethane concentration is therefore strongly indicative of atmospheric dispersion. Indeed, this characteristic of ethane is easily confirmed by plotting the diurnal and seasonal variation which shows highest concentrations at night and during the winter months i.e. when the atmosphere tends to be more stable. Because much of the information value of the surface meteorological variables relates indirectly to the atmospheric dispersion, it was posited that much of the information from surface measurements could be obtained from the ethane concentration. If this was found to be the case, it would, where measurements of ethane are available, enable six explanatory variables to be replaced by one within the air quality model, leading to a reduction of complexity and therefore improvements in interpretability, as well as reductions in computational demand. Additionally, it would eliminate the need to arbitrarily choose a meteorological monitoring site from which to collect the measurements, a choice that can, in the absence of discriminating information or detailed site knowledge, be highly subjective.

PM₁₀ concentrations at three different London monitoring sites, London Marylebone Road (MY1), Cromwell Road (KC2) and Camden Swiss Cottage (CD1) between 2000 and 2017 were modelled using random forest. For each site, three models were trained: one using surface meteorological data from NOAA (2016), one using ethane concentration measurements from MY1, and a control model, which contained no meteorological data at all. The input data for all other variables (the response i.e. PM₁₀ concentration, the background PM₁₀ concentration, traffic counts, and temporal variables) were identical for all three models. A normalised trend in PM₁₀ concentration was calculated using each model, as shown in Figure 3.2.




Figure 3.2: Comparison of the normalised trend in PM₁₀ concentration at (a) London Marylebone Road (MY1), (b) Cromwell Road (KC2), and (c) Camden Swiss Cottage (CD1) 2000-2017, calculated using models trained on NOAA surface meteorological measurement data 'noaa', hourly ethane concentration measured at MY1 ('ethane'), and on no meteorological data at all ('none').

For all three sites there was little observable difference in the normalised trends generated using models trained using meteorological data from the different sources. This is an unexpected result, but even more unexpected is the finding that neither of these trends differed significantly from the trend produced by the model trained on no meteorological data at all. This result implies that virtually all of the information contained within the meteorological data is redundant in explaining variation in PM_{10} concentrations, in the presence of the other predictors in the model. Since PM_{10} concentrations are known to be strongly influenced by the background concentration, it is likely that this variable provides the majority of the model's explanatory power. The strong inter-dependence of the predictors may also be a factor: meteorology and background concentrations are strongly inter-related, and therefore the information encoded within the meteorological data may, in the absence of such data, also be provided by the background variable. These hypotheses could be investigated in future work by comparing models trained without the background variable.

It is worth noting that the redundancy of the meteorological data has been observed only when modelling PM_{10} concentration. Other air pollutants of interest, such as NO_x and NO_2 , are more strongly influenced by local emissions, and models that predict these pollutants may therefore display different responses to the input of different meteorological data. Future work could repeat this analysis for these pollutants, to determine whether the same conclusions hold true.

3.3.2.2 Normalised Trend Estimation using Partial Dependence Plots

Partial dependence plots and the 'meteorological normalisation' method developed by Grange et al. (2018) rely on the same principle: removal the effect of all but a single predictor through the use of the trained model to predict the value of the outcome for each observation using many (or all) values of the other predictors. The average prediction for the observation represents the value of the outcome under 'average conditions' of all other predictors (more detailed descriptions of the partial dependence calculation and the random forest normalisation method are given in Section 3.1.4.2 and Section 3.1.5 respectively). The meteorological normalisation framework can be viewed as a specific case of the partial dependence function, where the predictor of interest is always the date (i.e. the long term trend term). Additionally, there are minor differences in the algorithms used to produce the outputs: the calculation of the partial dependence function involves calculating the average prediction given *all* values of the other predictors, whereas the meteorological normalisation algorithm draws a random

sample of user-specified size of predictor data with replacement (in this analysis, a sample size of 500 was used).

Another noteworthy difference between the two methods is their relative computational efficiency. The calculation of partial dependence plots in R (*pdp* package) is optimised, and can therefore be executed far more quickly than normalisation. It therefore offers the advantage of a faster execution speed, which is valuable when many normalised trends must be calculated many times (e.g. for multiple monitoring sites).

The similarity in the partial dependence and meteorological normalisation methods led to the conjecture that the two methods may generate similar outputs, but one may offer advantages in terms of execution speed over the other. PM_{10} concentrations at London Marylebone Road 2000-2017 were modelled using random forest, and the long term trend component was estimated using the normalised trend (from the meteorological normalisation method) and the partial dependence plot for the Unix date predictor. A comparison of the two trends is shown in Figure 3.3.




Figure 3.3: Comparison of the normalised trend and the partial dependence plot for the Unix date (long term trend trend) variable from the random forest model of PM_{10} concentration at Marylebone Road, London 2000-2017.

Figure 3.3 shows that the features of the normalised trend are represented accurately by the partial dependence plot. The partial dependence plot is smoother than the normalised trend, and lacks much of the shorter term variation, most likely due to the use of a greater number of observations in the calculation of the

average prediction for each observation. Thus, for tasks where conservation of computational resource is a crucial factor, the long term trend component could be represented by the partial dependence plot for the date variable, rather than the normalised trend.

However, the normalisation methodology offers advantages in terms of the flexibility of the method, and the level of detail it provides. For example, the meteorological normalisation method enables the user to filter the data from which the values of the predictors are sampled during normalisation. This allows the normalised trend to be calculated for specific conditions, rather than the 'average' conditions of the entire data set. In Chapter 4, for instance, this capability is leveraged to sample data from a single year, for each of the years in the data set in turn, in order to estimate the variation in pollutant concentration resulting from inter-annual meteorological variation.

In conclusion, the long term trend in air pollutant concentration can be represented either by the normalised trend, or by the date partial dependence plot. The choice of method for calculating the trend must be made on a case-by-case basis, according to the requirements of the analysis.

3.3.3 Long Term Normalised Trends in London

The normalised trends in NO_x , NO_2 and PM_{10} concentration at three sites in London between 2000 to 2017 were estimated using random forest models of hourly ambient monitoring data, as described in Section 3.2. These trends are shown in Figure 3.4.

CHAPTER 3. DEVELOPMENT AND APPLICATION OF RANDOM FOREST MODELS TO AIR POLLUTANT TIME SERIES

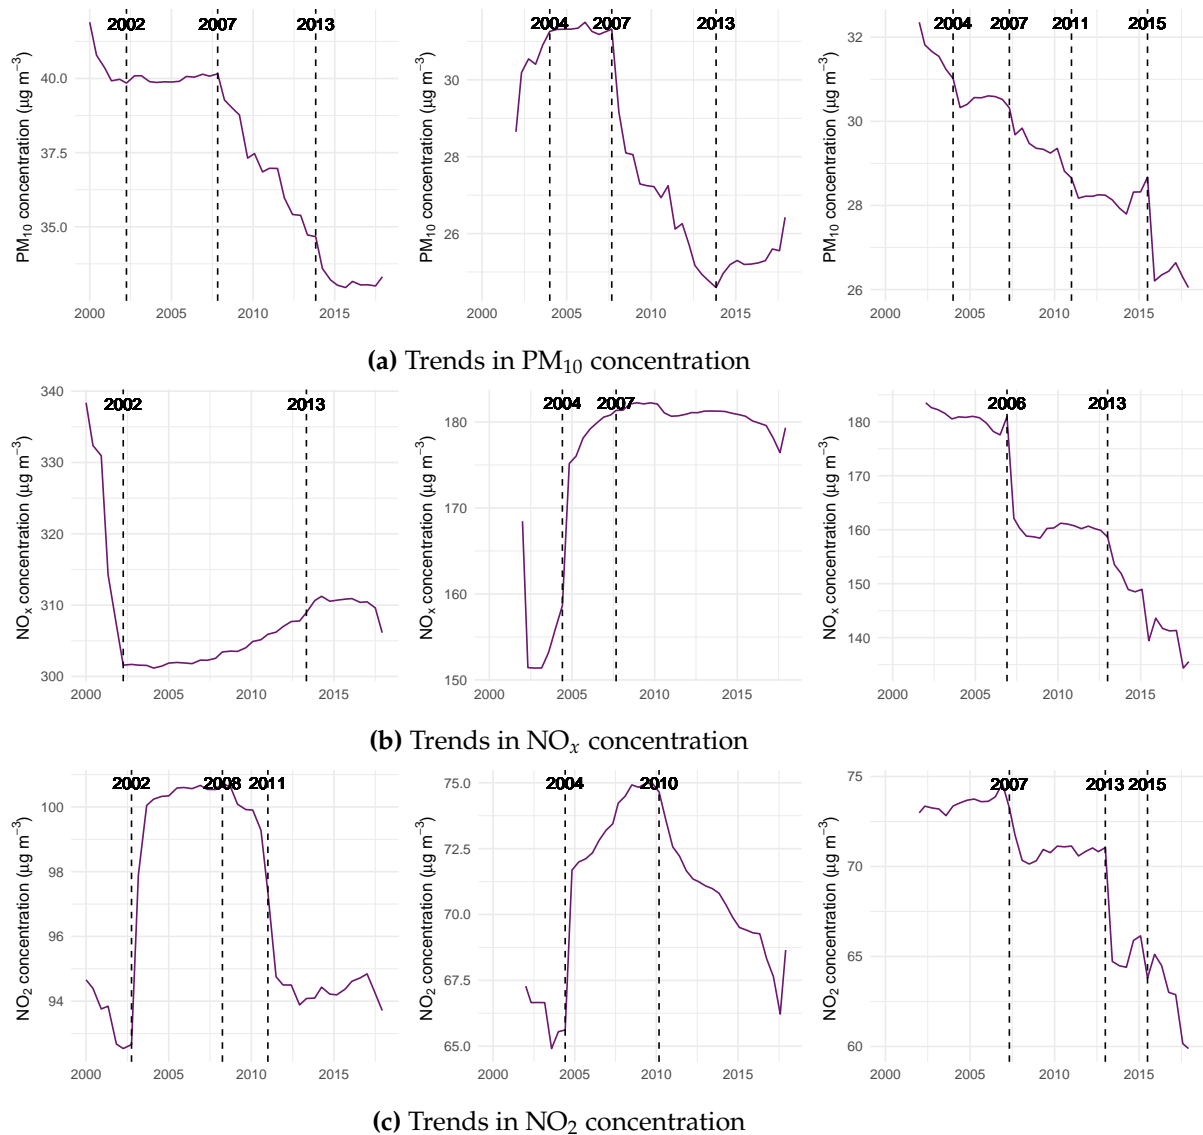


Figure 3.4: Long term trends in (a) PM₁₀, (b) NO_x, and (c) NO₂ concentration at, from left to right, Marylebone Road (MY1), Camden Kerbside (CA1), and Cromwell Road (KC2) in London between 2000 and 2017. The purple line is the partial dependence plot of the date variable calculated from the random forest model trained using monitoring data from the specified site. The vertical dashed lines indicate the break points.

Several common patterns can be observed in the long term normalised trends in Figure 3.4. All three sites exhibit a change point in 2007 where the PM₁₀ concentration begins to decrease monotonically. This date is ahead of the introduction of Phase 1 of the London LEZ that focussed on reducing PM emissions on 4th February 2008. The observed decline in PM₁₀ concentration could be driven by the gradual penetration of Euro 4/IV, and eventually Euro 5/V, vehicles into the fleet. PM₁₀ concentrations then level off, and then start to increase again at Camden Kerbside after 2013. It is possible that saturation of the vehicle fleet

with Euro 4-6 vehicles occurred at this point, and so further fleet turnover did not result in further decreases in emissions of PM. The reasons for the increase in PM₁₀ concentration after 2013 at Camden Kerbside are unknown, but may be the result of site-specific changes not included in the model (e.g. changes in road structure, congestion or the urban environment).

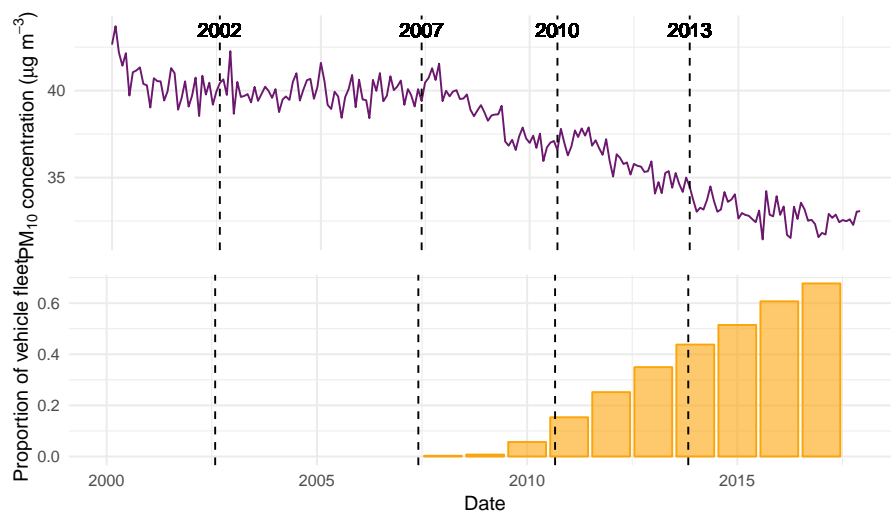
The long term trends in NO_x differ between all three sites. At Marylebone Road NO_x concentrations decrease sharply until 2002, followed by a gentle increase to 2013, then level off. At Camden Kerbside, NO_x concentrations increase until 2007, then plateau. At Cromwell Road, the trend is relatively constant except for step decreases in 2006/7 and a decrease after 2013. Other studies by [Grange et al. \(2017\)](#) and [Lang et al. \(2019\)](#) analysed the long term trends in ambient NO_x concentration in Europe, and consistently found that NO_x concentrations decreased monotonically between 2000 and 2017. [Lang et al. \(2019\)](#) also studied the long term trends in NO_x concentration in 18 individual European cities, concluding that the shape of the trend was remarkably consistent across a wide range of locations, once the effects of local variation were averaged out through aggregation of data from multiple monitoring sites. The marked difference in the shape of the trends at the Marylebone Road and Camden Kerbside sites from those expected based on literature findings suggests that the trends are strongly influenced by local effects, and could be interpreted through consideration of the history of the sites. It is likely therefore, to characterise the concentration changes in an area requires the consideration of many sites (as considered in Chapter 2), even if the meteorological variation can effectively be removed.

Marylebone Road and Camden Kerbside show similar trends in NO₂ concentration. Initially, the NO₂ concentration is observed to decrease until 2002/2004, before increasing between 2002/4 and 2010, and then decreasing. At Marylebone Road this is followed by a period of no change. This is consistent with the findings of [Grange et al. \(2017\)](#), who found that average NO₂ concentrations in Europe increased between 2000 and 2009, then decreased between 2010 and 2015. The initial increase in NO₂ concentration can be attributed to an increase in the number of diesel passenger vehicles over this period. Another driver of the observed increase could be increased primary emissions of NO₂ from vehicle exhaust in compliance with Euro 3 and 4 emission standards, brought about by the introduction of emission technologies such as DPF (diesel particulate filters) and DOC (diesel oxidation catalysts). These post-exhaust treatments deliberately oxidise NO to NO₂ in order to oxidise other combustion products, such as CO and particulate matter.


At Cromwell Road, the NO_2 concentration follows an overall pattern of decrease, with sharp step changes in 2007 and 2013 interspersed with period of little change. The step changes could be the result of changes in emissions resulting in the introduction of the emission technologies discussed previously, however the other monitoring sites' trends indicated that the changes in pollutant concentration due to these drivers typically occurred more gradually, as new vehicles penetrated the fleet. The sudden nature of the changes observed at Cromwell Road might suggest a different fleet composition, or an unusually fast rate of fleet turnover. More information would be necessary to draw firm conclusions.

It was suggested that many of the observed changes in air pollutant concentration were driven by changes in vehicle emission technology, and the gradual penetration of these newer vehicles into the fleet. To test this hypothesis, data on the UK vehicle fleet composition by emission class was compared to the normalised trends in PM_{10} , NO_x and NO_2 concentration at London roadside sites, as shown in Figure 3.5 for PM_{10} concentrations.


CHAPTER 3. DEVELOPMENT AND APPLICATION OF RANDOM FOREST MODELS TO AIR POLLUTANT TIME SERIES



(a) Marylebone Road (MY1)



(b) Camden Kerbside (CA1)



(c) Cromwell Road (KC2)

Figure 3.5: Comparison of the normalised trend in PM₁₀ concentration (purple line) at (a) Marylebone Road, (b) Camden Kerbside, and (c) Cromwell Road 2000-2017 with the proportion of the UK vehicle fleet composed of post-Euro 5 vehicles over the same period (yellow bars). The vertical dashed lines indicate the break points.

Figure 3.5 shows that break points occur around 2007-10 at all three sites, corresponding to the introduction of Euro 5 vehicles. Further break points occur in 2012, 2013 and 2015 as PM_{10} concentrations decrease further, and the vehicle fleet moves increasingly towards Euro 5 and 6 vehicles.

It should be noted that the trends in PM_{10} , NO_x and NO_2 concentration all exhibit too much variability for the results of the break point analysis to be statistically significant. There is therefore a risk of over-interpreting the results of the break point analysis. In reality, clear, sharp change points in the trends would not be expected, since turnover of the vehicle fleet, and therefore changes in the proportion of vehicles equipped with better emission technologies, are gradual and occur slowly over time. The subsequent reductions in roadside emissions of air pollutants would also be gradual, and therefore the concentrations at the roadside would also be expected to decrease gradually. This implies that in future work, break point analysis may not be the most appropriate method by which to analyse the changes in pollutant concentrations observed in the trends.

3.3.4 Intervention analysis of the London Low Emission Zone

The normalised trends in PM_{10} concentration were calculated for nine roadside London sites. The break points of these trends were calculated, and compared to the dates of the LEZ implementation phases. An intervention variable was included in the random forest model (taking values 0, 1, 2, 3 for the different phases of the LEZ). Normalised trends were generated with the indicator variable set to each of these values, and the trends were compared. An example of these normalised trends is shown for the Marylebone Road site in Figure 3.6.




Figure 3.6: Comparison of the normalised trends in PM₁₀ concentration at Marylebone Road, London 2008-2017 with the intervention indicator variable set to the values for Phases 0-3. These trends represent different counterfactual scenarios (i.e. the observed trend had each phase of the LEZ been in place for the entire time period).

In Figure 3.6, similar trends were observed for each counterfactual (i.e. for every phase): PM₁₀ concentrations decreased until 2015, then rose slightly and levelled off. However, the PM₁₀ concentration in the presence of the LEZ was lower than in its absence. For Phase 1, this difference was small, reflecting perhaps the relatively small amount of data associated with this phase, as well as the less ambitious restrictions it imposed. The PM₁₀ concentration associated with Phase 2 was lower than for Phase 1, and the PM₁₀ concentration for Phase 3 lower still. This is consistent with the expectation that the successive phases of the LEZ were responsible for successively lower emissions of PM₁₀ because fewer high-emitting vehicles entered the area encapsulated by the LEZ.

The average difference in PM₁₀ concentration between each combination of phases over the entire period of analysis (2010–17) for each monitoring site is shown in Figure 3.7. The error bars represent the uncertainty in the average. A paired t-test was applied to determine whether the differences in average PM₁₀ concentration estimated for different phases was statistically significant. The results of the paired t-test are shown in Figure 3.7 by the different colours. Blue indicates a significantly lower concentration in the higher phase, red indicates a significantly higher concentration in the higher phase, and green indicates that

the difference between the two phases was not statistically significant.

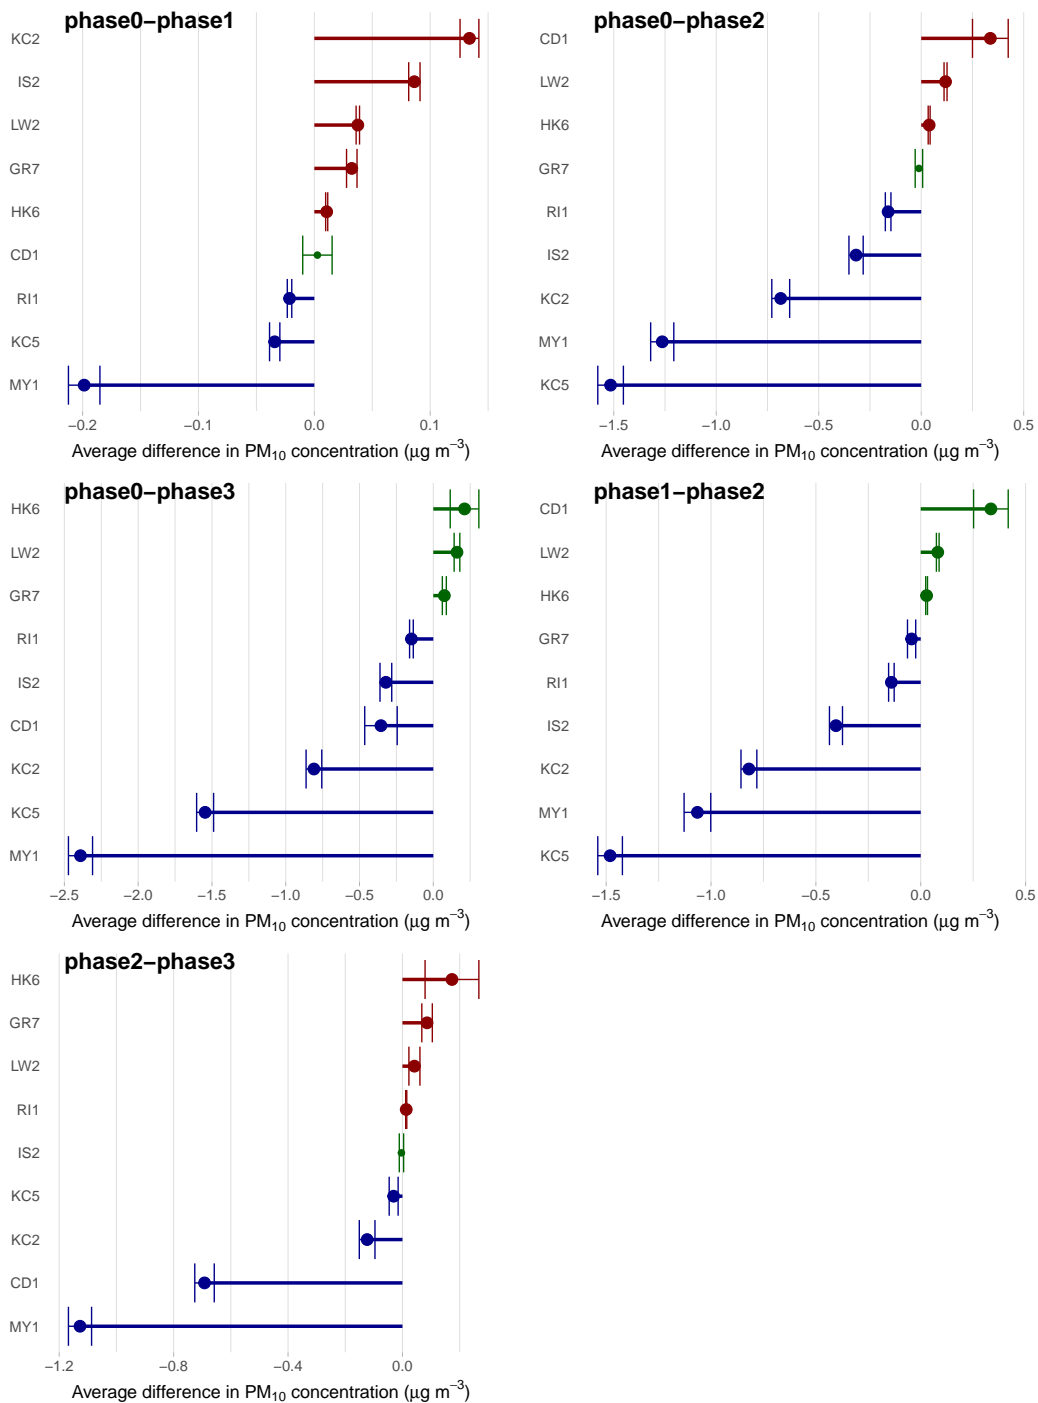


Figure 3.7: Summary of the results of the London LEZ intervention analysis (the effect of the London LEZ on ambient PM₁₀ concentrations) carried out using data from nine London monitoring sites. The segments indicate the average difference in PM₁₀ concentration between normalised trends for two different phases (values of the intervention indicator variable), with red indicating a higher PM₁₀ concentration for the higher phase, blue indicating a lower concentration, and green indicating no significant change (paired t-test).

Most sites showed significantly lower average PM_{10} concentrations in Phase 3 of the LEZ relative to the control scenario (no LEZ/Phase 0). The sites Hackney Old Street (HK6), Lewisham New Cross (LW2) and Greenwich Blackheath (GR7) showed slightly higher PM_{10} concentrations in Phase 3 of the LEZ, but these were not statistically significant. Similarly, most sites had lower PM_{10} concentrations for Phase 2 compared to Phase 1, with the exception of Camden (CD1), Lewisham New Cross (LW2) and Hackney Old Street (HK6), which showed statistically insignificant higher concentrations for Phase 2.

However, comparisons between other combinations of phases revealed a more ambiguous picture. The changes in average PM_{10} concentration between Phases 0 and 1, Phases 0 and 2, and Phases 2 and 3 showed higher concentrations for some sites and lower concentrations for others, with no clear consensus.

Overall, even the largest differences in average PM_{10} concentration between different scenarios were minor. The largest difference was observed for Marylebone Road (MY1), which showed that the average PM_{10} concentration in the Phase 3 scenario was $2.4 \mu\text{g m}^{-3}$ lower than in the no LEZ (Phase 0) scenario. Even the comparison between the no LEZ (Phase 0) scenario and the most restrictive LEZ scenario (Phase 3) showed that 3 out of the 9 sites studied, or 33% of the sites, showed no statistically significant difference in average PM_{10} concentration between the two scenarios. In summary, despite the sensitivity and proven accuracy of the technique, the evidence for whether the implementation of the LEZ has reduced air pollution in London was inconclusive. Any changes in PM_{10} concentration resulting from the London LEZ in the areas studied over this period were either small or insignificant.

3.4 Conclusions & Next Steps

3.4.1 Conclusions

This chapter has explored the use of random forest to produce meteorologically normalised ambient trends in considerable detail, including investigation of the method by its application to well-known changes (for example, changes in ambient black carbon concentrations at London Marylebone Road), comparison of the model performance and outputs under different modes of use (for example, with different sources of input meteorological data), and application of the method to intervention analyses (for example, analysis of the impact of the London LEZ on ambient air quality in London).

The method produced a normalised trend in black carbon concentration consistent with the expectations based on well known changes in vehicle emissions. This confirmed the effectiveness of the method in removing the obscuring influence of meteorology from the time series to reveal the comparatively small changes due to changes in emissions.

Similar normalised trends were produced using models trained using different meteorological data inputs, as well as a model trained using no meteorological data at all. This implies that for PM₁₀ at least, the information represented by the meteorological data is redundant with information provided by other predictors, such as background PM₁₀ concentration. The underlying reason for this behaviour is the dominant contribution from regional sources, which are well-captured by the use of urban background sites in the model. Alternatively it could be useful in the future to focus only on increments above background concentration. However, such an approach is prone to local influences at background sites and the generally very small concentration increment that exists between roadside and background sites.

The results of the investigation into the effectiveness of the London LEZ were not definitive. Different monitoring sites showed a range of different results, with the unexciting consensus view that later stages of the LEZ had probably resulted in small improvements in ambient air quality at most sites, however these proved to be only in the range of a few $\mu\text{g m}^{-3}$. Furthermore, while the majority of sites showed slightly, but significantly, lower concentrations of PM₁₀ in the presence of the LEZ, the differences at 33% of the studied sites proved not to be statistically significant. Further analysis may be able to reveal more information, or improve our confidence in the existing conclusions. Replacing the absolute concentrations with the increment above the background concentration may provide more definitive results and could form the focus of future work.

Bibliography

Bai, J., 1994. Least Squares Estimation of a Shift in Linear Processes. *Journal of Time Series Analysis* 15 (5), 453–472.

URL <http://doi.wiley.com/10.1111/j.1467-9892.1994.tb00204.x>

Department for Environment, Food and Rural Affairs (Defra), 2019. Automatic Urban and Rural Network (AURN).

URL <https://uk-air.defra.gov.uk/networks/network-info?view=aurn>

Department for Transport, 2010. GB Road Traffic Counts.

URL <https://data.gov.uk/dataset/208c0e7b-353f-4e2d-8b7a-1a7118467acc/gb-road-traffic-counts>

Dijkema, M. B. A., Van Der Zee, S. C., Brunekreef, B., Van Strien, R. T., 2008. Air quality effects of an urban highway speed limit reduction. *Atmospheric Environment* 42, 9098–9105.

URL <https://doi.org/10.1016/j.atmosenv.2008.09.039>

Donges, N., 2019. A Complete Guide to the Random Forest Algorithm.

URL <https://builtin.com/data-science/random-forest-algorithm>

Grange, S., 2018. rmweather: Tools to conduct meteorological normalisation on air quality data.

URL <https://github.com/skgrange/rmweather>

Grange, S. K., Carslaw, D. C., Lewis, A. C., Boleti, E., Hueglin, C., 2018. Random forest meteorological normalisation models for Swiss PM₁₀ trend analysis. *Atmospheric Chemistry and Physics Discussions*, 1–28.

URL <https://doi.org/10.5194/acp-18-6223-2018>

Grange, S. K., Lewis, A. C., Moller, S. J., Carslaw, D. C., 2017. Lower vehicular primary emissions of NO₂ in Europe than assumed in policy projections. *Nature Geoscience* 10 (12), 914–918.

URL <http://www.nature.com/articles/s41561-017-0009-0>

BIBLIOGRAPHY

Hoare, J., 2019. How is Variable Importance Calculated for a Random Forest? | Displayr.

URL <https://www.displayr.com/how-is-variable-importance-calculated-for-a-random>

James, G., Witten, D., Hastie, T., Tibshirani, R., 2013. An Introduction to Statistical Learning, 7th Edition. Vol. 103. Springer.

URL <http://www.springer.com/series/417><http://link.springer.com/10.1007/978-1-4614-7138-7>

Kings College London, 2019. London Air Quality Network.

URL <https://www.londonair.org.uk/LondonAir/Default.aspx>

Lang, P., 2018. rfmodels.

Lang, P. E., Carslaw, D. C., Moller, S. J., 2019. Analysis of long term roadside air pollution trends in European cities using ambient data from sparse monitoring networks. In: 23rd International Transport and Air Pollution Conference.

Lee, C., 2017. Feature Importance Measures for Tree Models — Part I.

URL <https://medium.com/the-artificial-impostor/feature-importance-measures-for-tree-models-part-i-47f187c1a2c3>

NOAA, 2016. Integrated Surface Database.

Ricardo Energy & Environment, 2019. Air Quality England.

URL <http://www.airqualityengland.co.uk/>

Sarica, A., Caresa, A., Quattrone, A., 2017. Random Forest Algorithm for the Classification of Neuroimaging Data in Alzheimer's Disease: A Systematic Review. *Frontiers of Aging Neuroscience* 9 (329).

URL <https://www.frontiersin.org/articles/10.3389/fnagi.2017.00329/full>

Swalin, A., 2018. Choosing the Right Metric for Evaluating Machine Learning Models — Part 1.

URL <https://medium.com/usf-msds/choosing-the-right-metric-for-machine-learning->

Williams, M., Minjares, R., 2016. A technical summary of Euro 6/VI vehicle emission standards. Tech. rep., ICCT.

URL www.theicct.org

BIBLIOGRAPHY

Wright, R., 2018. Interpreting Black-Box Machine Learning Models Using Partial Dependence and Individual Conditional Expectation Plots. Tech. rep., SAS Institute Inc.

URL <https://www.sas.com/content/dam/SAS/support/en/sas-global-forum-proceedings/2018/1950-2018.pdf>

4. Quantifying the Effect of Inter-annual Meteorological Variation on Pollutant Concentrations

4.1 Introduction

One of the main uses of air quality monitoring data is for monitoring compliance with EU limits on air pollutant concentrations. For some pollutants, these limits are defined by the annual mean metric, for example, the EU limit on the concentration of NO₂, a pollutant of considerable concern for regulators, is defined by two metrics: a 1-hour mean value of 200 µg m⁻³ not to be exceeded more than 18 times a year, and an annual mean value of 40 µg m⁻³ (Defra, 2019). It is well known, however, that meteorology exerts a powerful influence over ambient pollutant concentrations, and year-on-year variations in meteorology can exert considerable effects on the annual average concentration of air pollutants in any given year. This raises the question: if the value of the annual average concentration can depend on variations in meteorology throughout the year, how appropriate is it as a metric for measuring compliance? For example, suppose that the average concentration of NO₂ at a monitoring site complies with EU limits in one year, but exceeds the limits in the next. Is this elevation in ambient concentration the result of a genuine increase in ambient concentration driven by an increase in emissions, or is it a consequence of abnormal meteorology in the latter year resulting in extended pollution episodes? How much does normal inter-annual variation in meteorology affect whether or not a monitoring site is 'compliant' with regulatory limits from year-to-year?

While it is understood that inter-annual meteorological variation affects the annual mean concentrations of pollutants, the degree to which it does so is not well understood. In order to provide an answer to the afore-stated question, the effect of inter-annual meteorological variation on the annual average concentration should be quantified, to enable the confidence with which annual averages can be used in compliance monitoring to be ascertained.

In this chapter, the magnitude of the effect of meteorology on ambient concentrations of air pollutants is determined and quantified to apply an uncertainty to the annual mean concentration. Using this approach, an assessment is carried out for each monitoring site of whether the inter-annual variations in meteorology affect whether or not a site is compliant from year to year. The recent years in which the meteorology resulted in higher concentrations of pollutants, years in which it resulted in lower concentrations, and years in which it had little/no significant effect are identified. The implications this has in terms of identifying appropriate years for use as baseline years in air quality modelling are discussed. The specific meteorological conditions responsible for the variations in ambient concentration,

and the physical and chemical mechanisms that drive these relationships are described. Finally, the implications of these findings on the appropriateness of current metrics used for compliance monitoring within the EU are discussed.

4.1.1 Meteorology and Air Quality

An assessment of the effects of meteorology on ambient concentrations of air pollutants must depend on a knowledge of the relationships between meteorology and air quality, as well as the physical and chemical mechanisms that drive these relationships.

Jiang et al. (2014) noted that a number of studies investigating the effects of meteorology on ambient NO₂ concentrations had shown that stable (anticyclonic or high pressure) systems were characterised by higher concentrations of NO₂, while NO₂ concentrations were lower under unstable (cyclonic or low pressure) systems. Surface meteorological measurements provide information about the weather system by describing its effects:

Under stable, anticyclonic conditions, sinking air causes a subsistence inversion, which is characterised by clear skies, lower wind speeds, and lower relative humidity. In summer, these conditions will result in instability during the day marked by high temperatures, and instability during the night with cooling temperatures during which vapour condenses as mist/dew. This leads to frequent inversions. In winter the clear skies can result in very cold temperatures. Anticyclonic conditions result in low dispersion, where pollution is trapped at low levels, and can build up, causing pollution episodes, particularly in the winter (Wayne, 2000; AQEG, 2004).

In contrast, unstable cyclonic conditions are characterised by low pressure, higher wind speeds and lower temperatures. Cloudy skies and potentially higher levels of precipitation result in higher relative humidity. These are high dispersion conditions, under which pollution disperses more rapidly, leading to lower ambient concentrations (Wayne, 2000).

The presence of these weather systems can be identified from surface meteorological measurements. Periods of low wind speed, high temperature (in summer) or low temperature (in winter), low relative humidity, and high atmospheric pressure indicate the presence of a high pressure anticyclonic weather system. These systems are stable, resulting in low dispersion, and therefore might be expected to be observed in conjunction with elevated concentrations of primary pollutants.

In contrast, periods of high wind speed, lower temperature (in summer), high

relative humidity and lower atmospheric pressure indicate the presence of low pressure cyclonic weather systems, which are unstable, leading to high dispersion. Consequently, emissions of air pollutants (e.g. from vehicle exhausts at roadside) are quickly dispersed, leading to lower concentrations of air pollutants (Wayne, 2000).

Air pollution events, which can have a considerable influence over the annual average pollutant concentration in any given year, are often marked by anticyclonic weather systems.

4.2 Methods and Data

4.2.1 Description of Data

The data used 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, the Air Quality England database collected by Ricardo Energy & Environment, and the Scottish and Welsh air quality networks (SAQN, WAQN). They consisted of hourly observations of ambient NO_2 , NO_x and O_3 concentrations, measured at monitoring sites distributed throughout the UK. The concentrations of NO_x and NO_2 were measured using the European Commission reference method of chemiluminescence with molybdenum converter.

In this analysis, data from all monitoring sites that had co-located measurements of NO_x and NO_2 concentrations, and that met the data capture requirements (at least 80% data capture between 2008 and 2017) were included. This set of monitoring sites included 173 sites measuring NO_2 concentration, 161 sites measuring NO_x concentration, and 77 sites measuring O_3 concentration. Due to the requirement that all sites must measure NO_x and NO_2 , only a subset of the available O_3 monitoring sites were included in the analysis. Therefore it should be noted that the sites used may be biased towards locations where NO_x is more influential on O_3 , and the trends in O_3 and its relationship with meteorology may be influenced accordingly. Very different results and relationships may be highlighted in an analysis which included all sites measuring O_3 , which include remote rural sites. The geographical distribution of these sites is shown in Figure 4.1. More information about the sites is given in Appendix V.

CHAPTER 4. QUANTIFYING THE EFFECT OF INTER-ANNUAL METEOROLOGICAL VARIATION ON POLLUTANT CONCENTRATIONS



(a) NO₂ sites



(b) NO_x sites



(c) O₃ sites

Figure 4.1: Distribution of monitoring sites measuring (4.1a) NO₂, (4.1b) NO_x, and (4.1c) O₃ concentrations with at least 80% data capture between 2008 and 2017 in the UK.

Surface meteorological measurements were used as explanatory variables for predicting the ambient concentrations during random forest model training. These data were collected from the Integrated Surface Database (ISD) using the worldmet R package (NOAA, 2016; Carslaw, 2017). The variables included in the model were air temperature, air pressure, relative humidity, dew point, wind speed, wind direction and visibility. The meteorological data was obtained from the nearest meteorological site meeting the data capture requirements (at least 80% data capture between 2008 and 2017 for all seven of the aforementioned meteorological variables) to the monitoring site.

4.2.2 Modelling the Effect of Meteorology on Air Quality

The air pollutant concentration was modelled as a function of meteorology and temporal factors (Julian day, Unix date, hour of day) using random forest models.

Meteorological normalisation was applied to the concentration time series, using the trained model and the methodology described by Grange et al. (2018), with a modification. Grange et al. (2018) simulated the concentration under

'average conditions' by sampling a number of observations (e.g. 500) of all predictors except the date (which represents the long-term trend) from the data set, and predicting the value of the response (i.e. the pollutant concentration) for each sampled observation. The predicted data was then averaged by date, resulting in a data set of predicted values for n values of the date, which represent the values of the pollutant concentration under 'average' conditions. However, in this analysis, the normalisation process was carried out m times, where m is the number of years in the input data set. Each of the m normalised time series were generated according to the process above, but using only meteorological data from a single year. For example, the normalised data for the meteorological year 2008 was generated by repeatedly sampling meteorological data *from the set of meteorological data measured during the year 2008* in order to predict the response. In this way, the normalised time series of each meteorological year represented the pollutant concentration under the average meteorological conditions of that year.

Since the long-term (normalised) trend should be independent of inter-annual meteorology, we expect all m normalised trends to exhibit the same shape, however the absolute concentration of the trends should vary, depending on the meteorology of the year its data was sampled from.

By quantifying the variation in concentration among the m time series, the effect of the inter-annual meteorological variation over this period of time on the pollutant concentration was estimated.

4.2.3 Quantification of Variation

The effect of meteorology on ambient concentration was quantified by first calculating the average normalised concentration of the pollutant for each meteorological year (averaged over the normalised concentrations for the entire period of analysis). The uncertainty was then calculated as the range (maximum - minimum) of the distribution of annual mean normalised concentrations for the m meteorological years.

4.2.4 Visualisation of The Effect of Inter-annual Meteorological Variation on Air Quality Across the Monitoring Network

While the normalised trends can be directly visualised and compared for a single site, for the analysis of multiple monitoring sites across the UK, two methods were used to summarise and compare the normalised concentrations as a function of the meteorological year. These methods were heatmaps and cumulative

sums (CUSUM), a method for analysing network data developed by [Manly and Mackenzie \(2000\)](#).

4.2.4.1 Heatmaps

The heatmaps visualise the difference between the average predicted concentration in a given meteorological year, and the average predicted concentration from all meteorological years across the network of UK monitoring sites, as formalised in Equation 4.1.

$$f(x) = x_{ij} - \bar{x}_i \quad (4.1)$$

where x_{ij} is the mean predicted concentration of the normalised trend at site i calculated using data sampled from meteorological year j , and \bar{x}_i is the mean predicted concentration of the normalised trends across all meteorological years at site i .

For the purposes of comparability between sites with different absolute concentrations of the pollutant, the values on the colour scale were standardised as shown in Equation 4.2, where x_i is the observation, \bar{x} is the mean of the observations for a given site, and s is the standard deviation of the observations at a site.

$$z_i = \frac{x_i - \bar{x}}{s} \quad (4.2)$$

4.2.4.2 Cumulative Sums

Cumulative sum (CUSUM) plots are a useful tool for showing how measured concentrations deviate from the ‘business as usual’ scenario (i.e. concentration under ‘average’ meteorological conditions). The accumulated difference between a variable and ‘business as usual’ is plotted. In this way, it indicates change-points even when those changes are small. If the CUSUM is zero, the measured concentration has not deviated from ‘business as usual’ ([Carslaw, 2020](#)).

In this case, the CUSUM is the cumulative sum of the difference between the predicted concentration of the normalised trends for a given meteorological year and the average predicted NO₂ concentration across all meteorological years, for all UK sites. The CUSUM approach used in the current context considers a *network of monitoring sites* with the aim of showing both whether a particular meteorological year differs significantly from average conditions and the behaviour of individual

site responses. The cumulative sum was calculated as shown in Equation 4.3.

$$S_{ij} = (x_{1j} - \bar{x}_1) + (x_{2j} - \bar{x}_2) + \dots + (x_{ij} - \bar{x}_i) \quad (4.3)$$

where S_{ij} is the CUSUM at site i for meteorological year j , x_{ij} is the mean normalised concentration at site i for meteorological year j , and \bar{x}_i is the mean normalised concentration across all meteorological years at site i . The CUSUM plot visualises S_{ij} as a function of i (the site).

A randomisation process was used to calculate the statistical significance of the difference between the concentration in a given meteorological year and the average concentration across all meteorological years, over the monitoring network. This involves comparing the observed CUSUM plots to the CUSUMs predicted by the null hypothesis (that there is no difference in concentration between different meteorological years). A large number of CUSUMs are generated, where the observations from each site are randomly permuted. The maximum and minimum values of S_{ij} for each of the sets of permutations are plotted as an envelope on the CUSUM plot, and represent the bounds for which the null hypothesis could hold true (i.e. differences from 0 could be due to random variation within real data). Randomisation is also used to derive p-values for the statistical significance of the difference between the concentration in a given meteorological year and the average concentration over all met years. The detail of this calculation is described in [Manly and Mackenzie \(2000\)](#) and [Manly and Mackenzie \(2003\)](#).

CUSUM plots were calculated for all UK sites, based on four different site orderings. Applying different ways of ordering the monitoring sites enables more information about the variation in the effect of meteorology on pollutant concentration across the monitoring sites to be ascertained. While ordering the sites by latitude and longitude indicated the north-south and east-west spatial variation, ordering by the average pollutant concentration revealed the variation by the level of pollution at the monitoring site, and ordering by the average NO_2/NO_x ratio of the monitoring site showed how the effect of meteorology varied based on the contribution of the traffic source to air pollution at the site.¹

In terms of the interpretation of the CUSUM plots, a positive slope indicates that, on average, the concentration in a given meteorological year is higher than the average concentration over all meteorological years. A negative slope indicates

¹The average NO_2/NO_x ratio at a monitoring site provides a proxy for the degree of influence of the traffic source on air quality at the monitoring site. Sites with lower NO_2/NO_x ratios indicate a stronger influence from traffic sources (perhaps due to being close to a busier road), while sites with lower NO_2/NO_x ratios indicate sites where air quality is more representative of the background.

that, on average, the concentration in a given meteorological year is lower than the average concentration over all meteorological years. A change in the gradient of the CUSUM plot indicates a change in the magnitude of the deviation in concentration for that meteorological year between two groups of sites.

The average deviation from the mean concentration, in a given meteorological year, can be calculated as the slope of the CUSUM plot in that year. The p-value for a CUSUM plot indicates the statistical significance of the deviation from the average concentration in that year: i.e. the probability that the observed difference from the average concentration is due to random chance. The overall p-value (for all CUSUM plots) indicates whether the CUSUM plots at all times between them differ from the average i.e. the probability that the observed differences overall differ from the average.

4.3 Results and Discussion

Random forest models were trained to predict NO_2 , NO_x and O_3 concentrations as a function of meteorology between 2008 and 2017 for each UK sites measuring the pollutant with sufficient data capture over this period.

The average performances of the models are given in Table 4.1.

Table 4.1: Random forest model performance metrics (averaged over all monitoring sites). The root mean squared errors (RMSE) are given in $\mu\text{g m}^{-3}$.


Pollutant	R-squared	OOB RMSE	Test RMSE
NO_2	0.6	13.6	8.2
NO_x	0.6	49.4	30.0
O_3	0.7	11.3	6.7

For each site, the normalised concentration for each meteorological year 2008–2017 was predicted using the trained model and meteorological data from the given meteorological year. As an example, Figure 4.2 shows the normalised trends in NO_2 , NO_x and O_3 concentration for each meteorological year for the Marylebone Road site (left-hand plot), as well as the variation in the absolute concentration of the trend in different meteorological years over this period (right-hand plot). The variation is shown as an envelope representing the maximum and minimum values of the trends over all meteorological years.


Figure 4.2 indicates that meteorology exerts a large influence over the ambient concentrations of all three pollutants. It would, therefore, be valuable to quantify the variation in ambient concentration that can be attributed to variations in

CHAPTER 4. QUANTIFYING THE EFFECT OF INTER-ANNUAL METEOROLOGICAL VARIATION ON POLLUTANT CONCENTRATIONS


meteorology, to avoid over-confidence in the interpretation of long term trends, or mis-attribution of this variation to other factors.



(a) NO₂



(b) NO_x



(c) O₃

Figure 4.2: Meteorologically normalised time series for (4.2a) NO₂, (4.2b) NO_x, and (4.2c) O₃ concentration 2008-2017 at Marylebone Road, for each meteorological year. The plot on the left shows the trends for each meteorological year 2008-2017, and the plot on the right shows the range resulting from meteorological variation (the range of predicted concentrations over all meteorological years) as an envelope around the average normalised trend.

4.3.1 Quantifying the Effect of Inter-annual Meteorological Variation

The effect of inter-annual variations in meteorology between 2008 and 2017 on the annual mean concentrations of NO_x, NO₂ and O₃ was quantified to provide some measure of the range in the mean concentration. The range was calculated as the range (maximum – minimum) of the concentration values predicted using data from the meteorological years 2008–2017.

The range in the value of the average concentration due to variations in meteorology for each pollutant is visualised as error bars on the data points representing the annual mean concentration. This was done for each monitoring site individually, as shown in Figures 4.3, 4.6, and 4.8, as well as averaged over all the monitoring sites in the networks for each pollutant, as shown in Figures 4.5, 4.7 and 4.9.

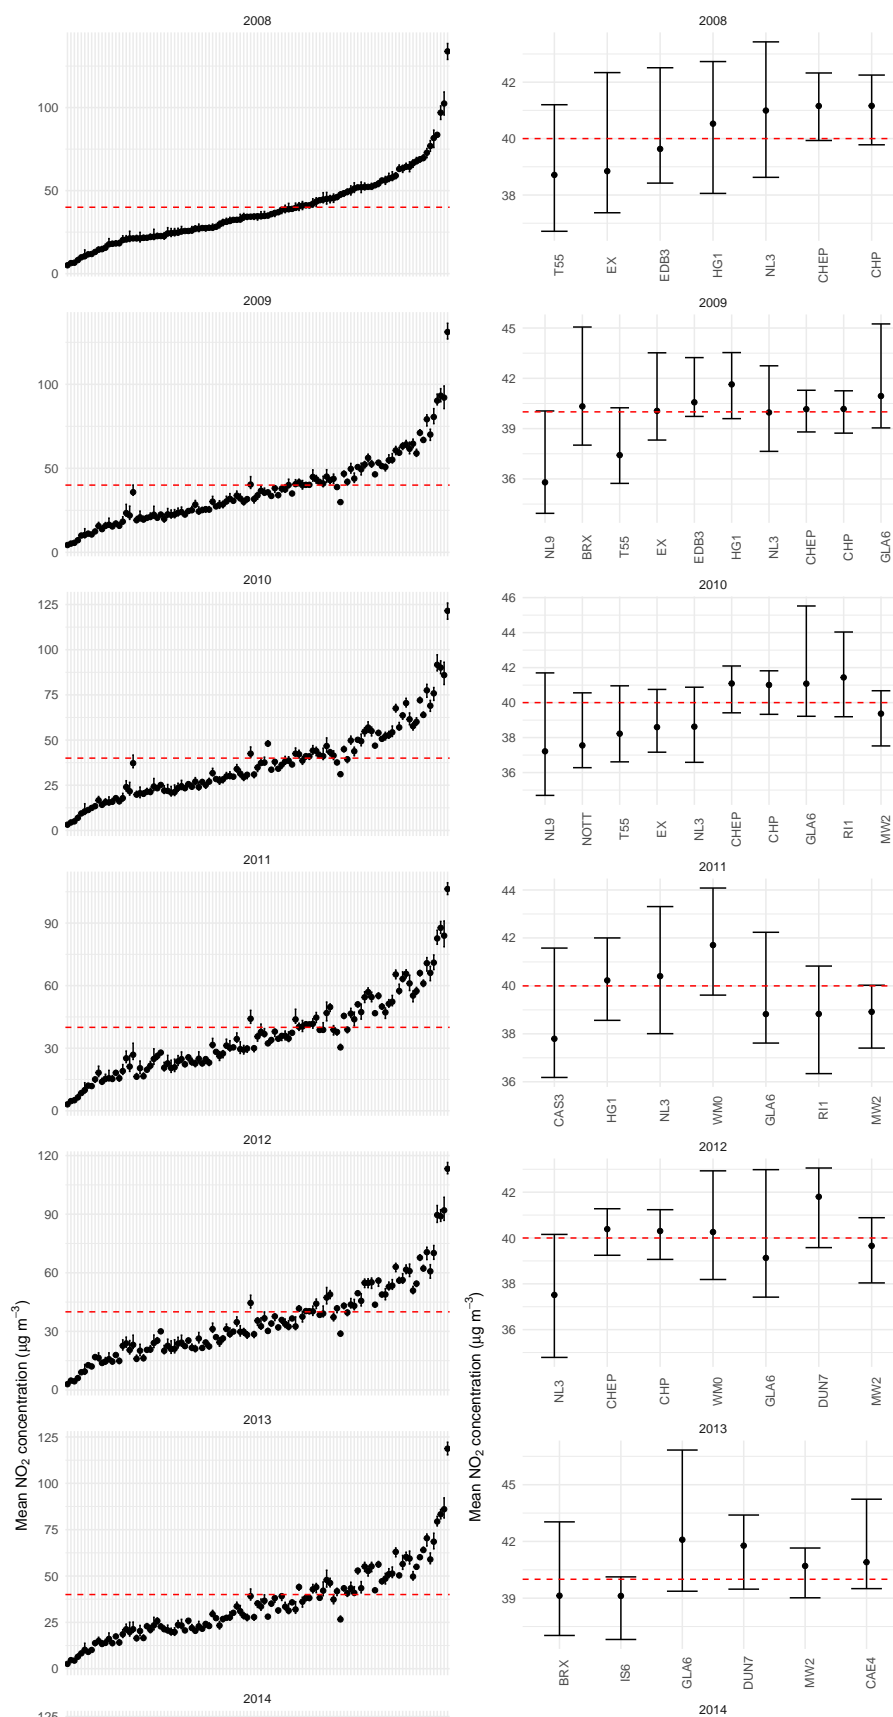
4.3.1.1 NO₂

Each row of Figure 4.3 represents the mean concentration and range due to inter-annual meteorological variation at each monitoring site in a given year. The dashed line marks the EU annual average limit value for NO₂ concentrations of

CHAPTER 4. QUANTIFYING THE EFFECT OF INTER-ANNUAL METEOROLOGICAL VARIATION ON POLLUTANT CONCENTRATIONS

$40 \mu\text{g m}^{-3}$. The plot on the left shows the data for all monitoring sites in the UK network, while the plot on the right is a zoomed in view that shows only those sites where the EU annual mean limit value for NO_2 concentrations lies within the range due to meteorological variation (subsequently referred to as 'marginal' sites). In other words, the right-hand plot focuses on the sites where the effects of meteorological variation could affect whether or not the site is compliant with the EU limits in that year.

CHAPTER 4. QUANTIFYING THE EFFECT OF INTER-ANNUAL METEOROLOGICAL VARIATION ON POLLUTANT CONCENTRATIONS



CHAPTER 4. QUANTIFYING THE EFFECT OF INTER-ANNUAL METEOROLOGICAL VARIATION ON POLLUTANT CONCENTRATIONS

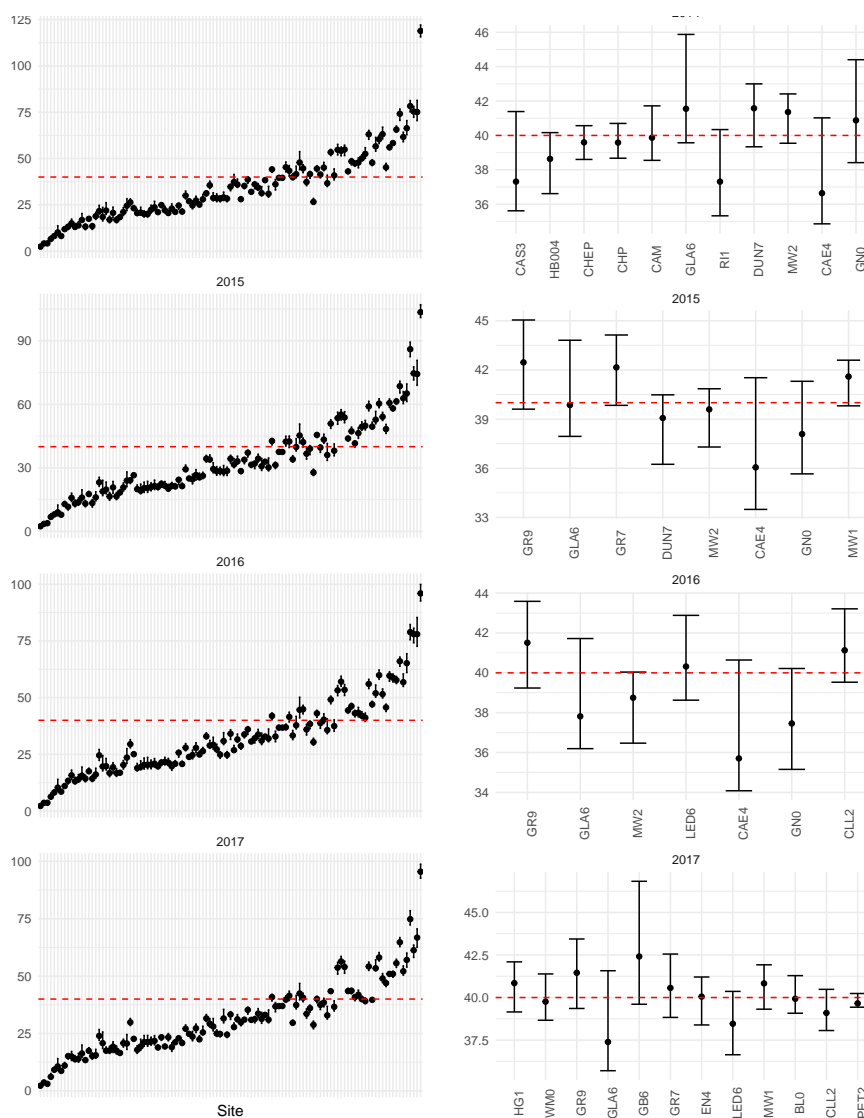


Figure 4.3: Range in the annual average NO₂ concentration at all UK sites in each year due to meteorology. The data points show the mean NO₂ concentration in that year at an individual monitoring site, while the error bars represent the range in the mean value resulting from meteorological variation. The range was calculated as the range in the predicted concentration for that year over normalised data from all meteorological years. The plot on the right is a zoomed in view of the plot of the left, showing the ‘marginal’ sites (those whose compliance with the EU limit values falls within the range due to meteorological variation) in more detail.

Over the entire period analysed (2008–2017), there were 44 ‘marginal’ monitoring sites for which the EU limit value for annual mean NO₂ concentration lay within the range resulting from meteorological variation in at least one year. The geographical distribution of these marginal sites is shown in Figure 4.4. Since a total of 173 sites were included in the analysis, this represents 25% of the total site

CHAPTER 4. QUANTIFYING THE EFFECT OF INTER-ANNUAL METEOROLOGICAL VARIATION ON POLLUTANT CONCENTRATIONS

coverage. The average number of sites which exceeded the EU limit value in a given year between 2008 and 2017 was 54.8. The average number of sites which were marginal in a given year was 12.1, therefore the average number of marginal sites as a proportion of the exceeding sites in any given year was 22%.



Figure 4.4: Spatial distribution of the sites which are ‘marginal’ for NO₂ (i.e. their compliance with EU limit values falls within the range (min to max) due to inter-annual meteorological variation in at least one year between 2008 and 2017).

Figure 4.5 shows the range (min to max) in the annual mean NO₂ concentration due to meteorological variation averaged over all UK sites. This range was calculated as the range of concentration values predicted using data from the meteorological years 2008–2017. On average, the annual range in the NO₂ concentration due to inter-annual meteorological variation was $2.9 \mu\text{g m}^{-3}$ (8.2%).





Figure 4.5: Annual average NO₂ concentrations, and range (min to max) due to meteorological variation (represented by the error bars), averaged across all UK sites over the period 2008 to 2017.

4.3.1.2 NO_x

The previous analysis of the range in annual concentration due to meteorological variation was repeated for the NO_x concentration data. As for Figure 4.3, each row of Figure 4.6 represents the mean concentration and range due to inter-annual meteorological variation at each monitoring site in a given year. However, the EU does not enforce limits on the NO_x concentration using the annual average metric, as it does with NO₂ concentration, therefore the limit values are not shown.

CHAPTER 4. QUANTIFYING THE EFFECT OF INTER-ANNUAL METEOROLOGICAL VARIATION ON POLLUTANT CONCENTRATIONS






Figure 4.6: Range in the annual average NO_x concentration at all UK sites in each year due to meteorology. The data points show the mean NO_x concentration in that year at an individual monitoring site, while the error bars represent the range (min to max) in the mean value resulting from meteorological variation. The range was calculated as the range in the predicted concentration for that year over normalised data from all meteorological years.

The range in NO_x concentrations resulting from meteorological variations, averaged over all UK sites, is shown in Figure 4.7.




Figure 4.7: Annual average NO_x concentrations, and ranges due to meteorological variation (represented by the error bars), averaged across all UK sites over the period 2008 to 2017.


On average, the range in the NO_x concentration due to inter-annual meteorological variation over the period 2008-2017 was $9.93 \mu\text{g m}^{-3}$ (12.6%).

4.3.1.3 O₃

The previous analysis of the range in annual concentration due to meteorological variation was repeated for the O₃ concentration data. However, the EU limits on ozone concentration are measured as an 8-hour rolling average, rather than the annual average concentration, therefore for ozone, the eight-hour rolling mean, and the associated range in the rolling mean resulting from meteorological variation was calculated, rather than the average concentration. Each data point in Figure 4.8 represents the annual average of the 8-hour rolling mean concentration in the specified year, while the error bars represent the range in the annual average of the 8-hour rolling mean due to inter-annual meteorological variation at each monitoring site in the given year. The EU Directive specifies that the daily maximum 8-hour rolling mean concentration must not exceed $120 \mu\text{g m}^{-3}$ for more than 25 days averaged over a 3 year period. Since this analysis considered the effects of meteorological variation on the *annual average* of the 8-hour rolling mean, rather than the *daily maximum* of the 8-hour rolling mean, the effect of meteorological variation on site compliance according to the EU Directive was not directly assessed, however the large magnitude of the percentage range implies that it may have a considerable effect. Further work might quantify the effect of interannual meteorological variability on the daily maximum of the 8-hour rolling

mean concentration in order to directly evaluate its effect on site compliance with the EU Directive. By examining this metric, more detail regarding the daily influence of meteorology on ozone concentration would be provided - and therefore the results might be quite different- however it would likely be necessary to consider each monitoring site one-by-one and therefore comparing the results across the entire network would be more challenging.

CHAPTER 4. QUANTIFYING THE EFFECT OF INTER-ANNUAL METEOROLOGICAL VARIATION ON POLLUTANT CONCENTRATIONS



CHAPTER 4. QUANTIFYING THE EFFECT OF INTER-ANNUAL METEOROLOGICAL VARIATION ON POLLUTANT CONCENTRATIONS

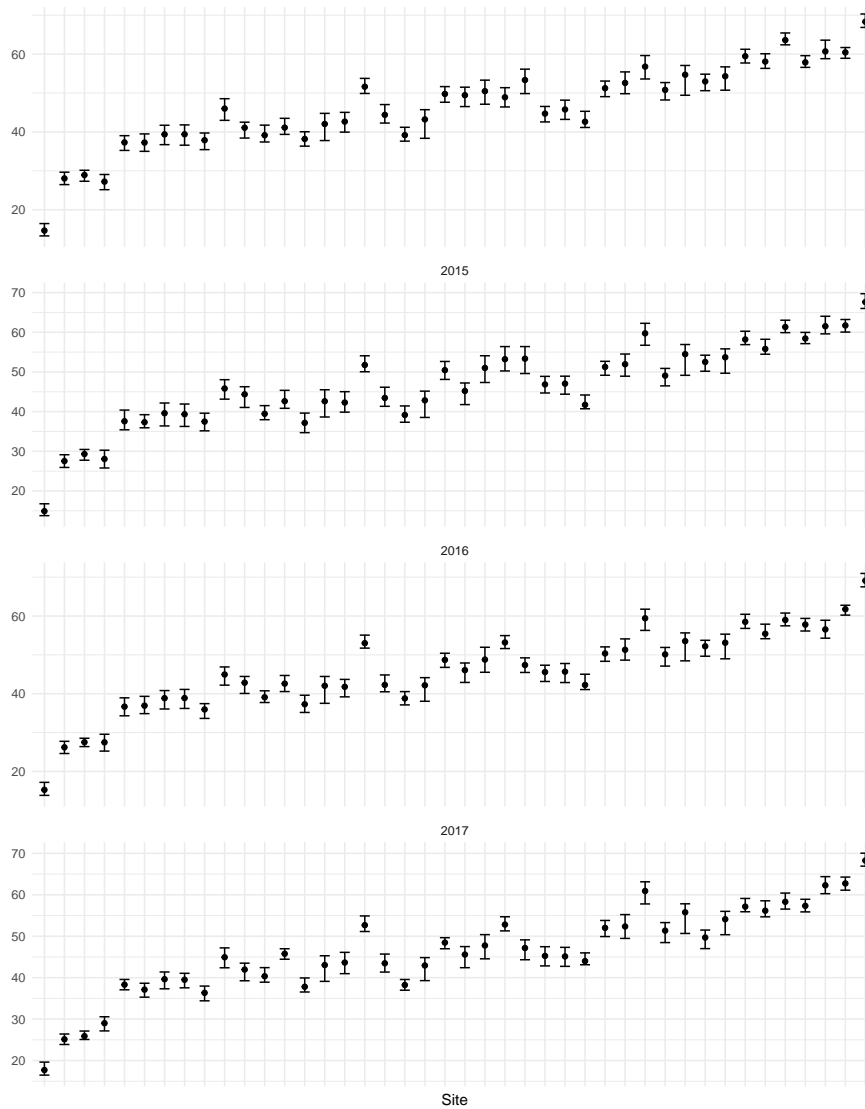


Figure 4.8: Range in the annual average O₃ concentration at all UK sites in each year due to meteorology. The data points show the mean O₃ concentration in that year at an individual monitoring site, while the error bars represent the range of the mean value resulting from meteorological variation. The range was calculated as the range in the predicted concentration for that year over normalised data from all meteorological years.

The range in O₃ concentrations resulting from meteorological variations, averaged over all UK sites, is shown in Figure 4.9.




Figure 4.9: Annual average O₃ concentrations, and ranges due to meteorological variation (represented by the error bars), averaged across all UK sites over the period 2008 to 2017.


On average, the range in the O₃ concentration due to inter-annual meteorological variation over the period 2008-2017 was 3.3 µg m⁻³ (7.5%). This variation is similar to that for NO₂ (2.9 µg m⁻³), which might be expected from the perspective of the inter-conversion between NO₂ and O₃ (See Section 1.1.2 for more detail).

4.3.2 The Effect of Inter-annual Meteorological Variation on Air Quality across the UK Monitoring Network


To evaluate the by-site effect of meteorology on annual average concentrations across the UK in more detail, the heatmap and CUSUM methodologies described in Section 4.2.4 were applied to the normalised NO_x, NO₂ and O₃ concentrations from each UK site in the monitoring network.

The heatmaps in Figure 4.10 are ordered by the mean concentration of the pollutant at each monitoring site.


CHAPTER 4. QUANTIFYING THE EFFECT OF INTER-ANNUAL METEOROLOGICAL VARIATION ON POLLUTANT CONCENTRATIONS



(a) NO_2



(b) NO_x



(c) O₃

Figure 4.10: Heatmaps showing the difference between the mean concentration in each meteorological year and the average concentration over all meteorological years for (4.10a) NO₂, (4.10b) NO_x, and (4.10c) O₃, for each UK monitoring site, as predicted by the random forest model.

A visual inspection of Figures 4.10a and 4.10b clearly indicates that 2010 stands out as a particularly ‘bad’ year for NO_x and NO₂, with the concentrations of these pollutants considerably elevated due to meteorology at almost all monitoring sites. Other (although less obvious) years with meteorology that was responsible for elevated NO₂ and NO_x concentrations at many sites are 2012, 2013 and 2016. On the other hand, the meteorology in 2015 was responsible for much lower-than-average concentrations of these pollutants at most sites, as, to a lesser extent, was the meteorology of 2008, 2014 and 2017.

In almost direct contrast, Figure 4.10c shows that the meteorology in 2010 was responsible for particularly *low* concentrations of O₃ at most sites, with less extreme reductions in concentration driven by meteorology in 2009, 2012, 2016 and 2017, while 2015 was characterised by meteorology that resulted in strongly elevated O₃ concentrations at most sites, with less extreme elevations in concentration also observed for 2008, 2011, 2013 and 2014. This behaviour

is expected based on the chemistry involved e.g. increased NO concentrations would tend to reduce O₃ (see reaction given by Equation 1.4).

Generally, there was good consensus among the monitoring sites regarding whether the meteorology in a given year was responsibly for higher or lower than average concentrations of each pollutant. However, there were some notable exceptions.

For example, at Hull Freetown (HUL2), meteorology caused elevated concentrations of NO₂ in 2009, 2011 and 2015, and lowered concentrations of NO_x in 2010 and 2016, in disagreement with the majority of sites.

Another notable site was Cambridge Roadside (CAM), which exhibits a distinct relationship between both NO₂ and NO_x concentrations and meteorology, with elevations of NO₂ and NO_x in 2011, 2013, 2015 and 2017, and lowered NO_x and NO₂ in 2009 and 2010.

Several of the most polluted sites (e.g. City of London Walbrook Wharf, Marylebone Road and Ealing Hanger Lane Gyratory) also differed strongly from the majority of sites in the effects of annual meteorology on NO_x and NO₂ concentration. This may be due to the strong influence of traffic emissions on concentrations at these sites (meaning that meteorology is less predictive of concentrations), as well as the type of environment. Often, these highly polluted sites are located at the roadside in highly built-up urban areas, such as street canyons, where the meteorology (in particular wind speed and wind direction) can be highly localised and complex, and possibly not well described by the regional meteorological information measured at the meteorological sites, which are often located at open locations (e.g. airports and airfields) some distance away.

Other sites with unusual effects of annual meteorology on NO_x or NO₂ concentrations include Islington Holloway Road (IS2), Cromwell Road (KC2), London Hillingdon (HIL), Hillingdon Keats Way (HI0), London Hillingdon Oxford Avenue (HI3), Chepstow A48 (CHEP, CHP), Chichester A27 Chichester Bypass (CI1), Aberdeen (ABD), High Muffles in Yorkshire (HM). Site information and geographical locations of all of the monitoring sites mentioned above are shown in Figure 4.11.

CHAPTER 4. QUANTIFYING THE EFFECT OF INTER-ANNUAL METEOROLOGICAL VARIATION ON POLLUTANT CONCENTRATIONS



Figure 4.11: Distribution of monitoring sites with unusual NO_x and NO_2 heatmap results (unusual variation in NO_x and NO_2 concentration as a result of inter-annual meteorological variation).

Table 4.2: Number of sites with unusual heatmap results for NO_x and NO_2 , by region and site type.

Region	Site type	Frequency
London	Urban traffic	5
London	Urban background	3
Eastern	Urban traffic	1
North East Scotland	Urban background	1
South East	Urban traffic	1
South Wales	Unknown	1
South Wales	Urban traffic	1
Yorkshire & Humberside	Rural background	1
Yorkshire & Humberside	Urban background	1

Table 4.2 summarises the information in Figure 4.11: it shows the number of monitoring sites, by region and site type, that show unusual NO_x and NO_2 heatmap results (i.e. which exhibit variations in NO_x and NO_2 concentration in response to inter-annual variations in meteorology that appear considerably different from most other UK sites). It can be seen from Table 4.2 and Figure 4.11 that the majority of the sites are located in London and the south-east. This

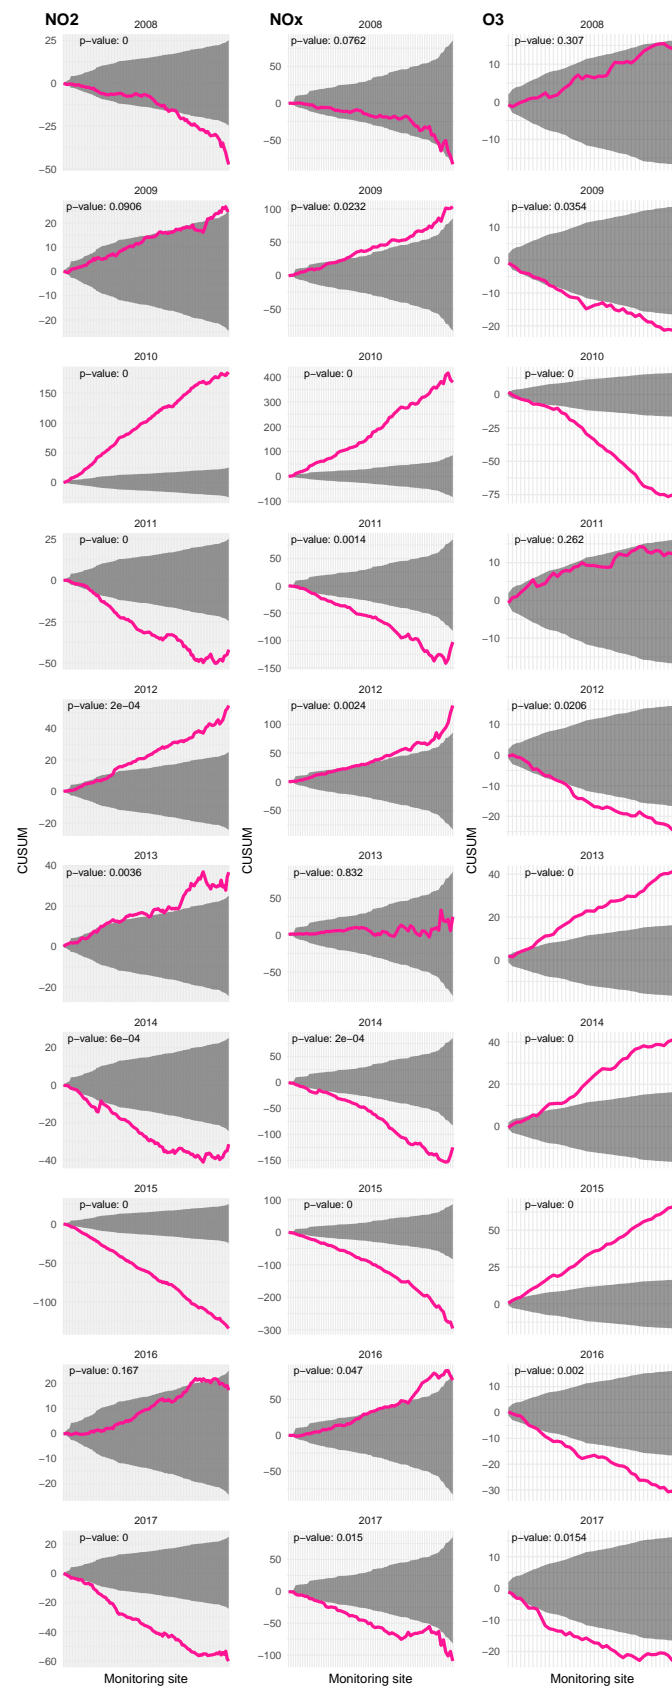
may be due to this region experiencing different meteorology from the rest of the UK, or, more likely, a greater proportion of sites located in street canyons or built-up areas with hyper-local meteorological conditions that are not well-represented by the regional meteorology. On the other hand, it may simply be due to the heterogeneity of monitoring site distribution in the UK: most of the total monitoring sites are located in London and the south-east, therefore most of the sites with unusual results are likely to be there too. There is no clear dependence on site type, with the proportion of traffic, urban background and rural sites roughly similar to their overall proportions among the total number of monitoring sites.

Figure 4.12 shows the CUSUMs of the difference in NO_x , NO_2 and O_3 annual concentration between each meteorological year and the average over all meteorological years for the meteorological years 2008–2017, with the sites ordered by mean concentration, NO_2/NO_x ratio, latitude and longitude. Each CUSUM plot is annotated with the p-value calculated using the randomisation process, which represents the statistical significance of the difference from the average concentration due to annual meteorology in that year (more precisely, it represents the probability that the observed difference from the average in that year is due to random chance, rather than the effects of meteorology).

The results shown by the CUSUM plots in Figure 4.12 match those identified earlier in the heatmaps (Figure 4.10). The CUSUM plots indicate that the most statistically significant deviation from normal meteorology (the null hypothesis) occurred in 2010 - a year in which concentrations of NO_x and NO_2 were abnormally high, and concentrations of O_3 abnormally low. Significantly lower-than-normal concentrations of NO_x and NO_2 and higher concentrations of O_3 were also seen in 2012, 2013, 2016 and 2009. The second-greatest deviation from the null hypothesis (i.e. 'normal' meteorology), was seen for 2015, which, in contrast to 2010, had higher than average concentrations of NO_x and NO_2 , and lower concentrations of O_3 . Similar, but less extreme, deviations were shown for 2014, 2017, 2011 and 2008.

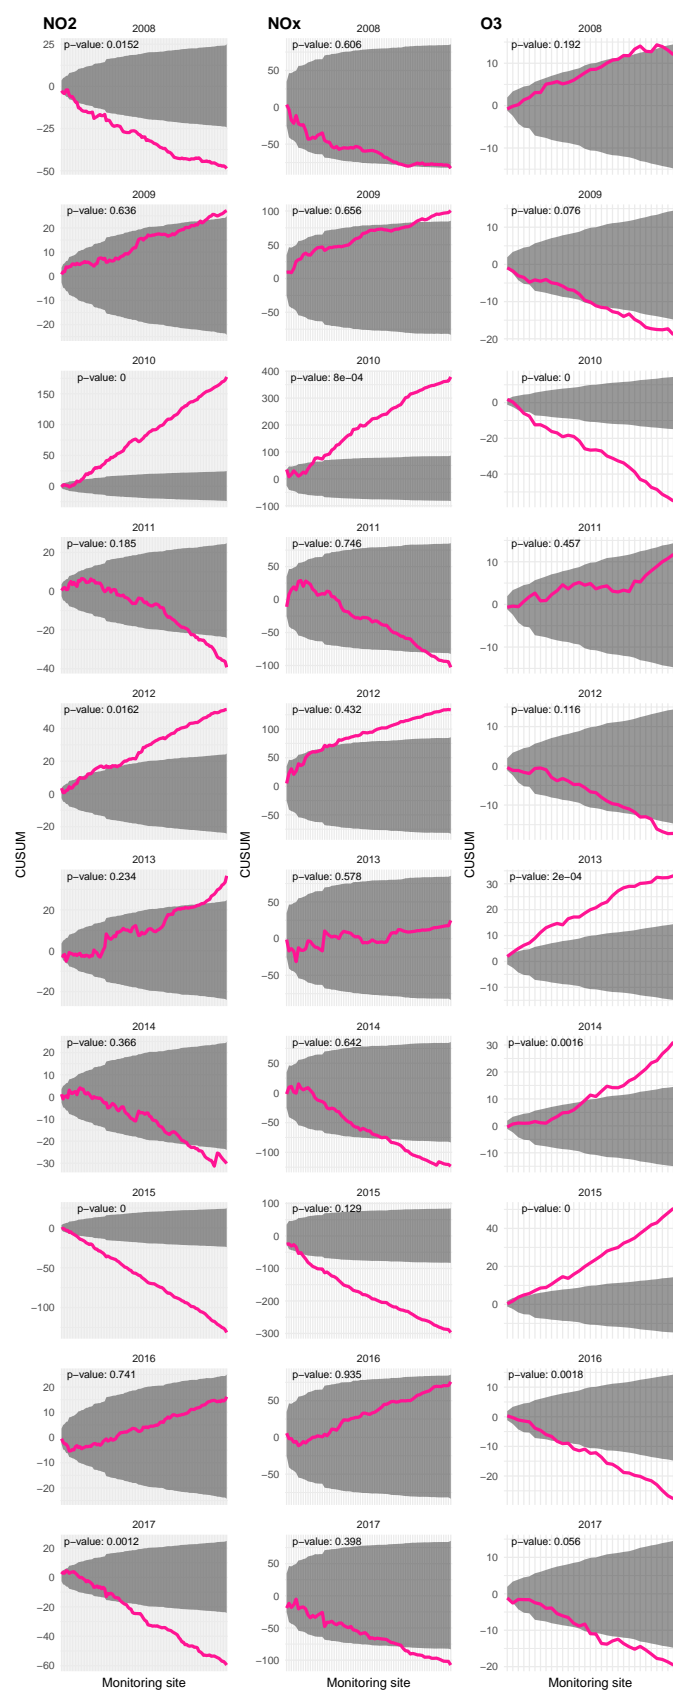
In general, the CUSUM plots for NO_2 and NO_x tend to mirror each other fairly closely, and will therefore generally be discussed together. The CUSUM plots for O_3 are often distinct (presumably because the concentrations of NO_x and O_3 are less strongly correlated, as the relationship is complicated by secondary chemistry).

CHAPTER 4. QUANTIFYING THE EFFECT OF INTER-ANNUAL METEOROLOGICAL VARIATION ON POLLUTANT CONCENTRATIONS



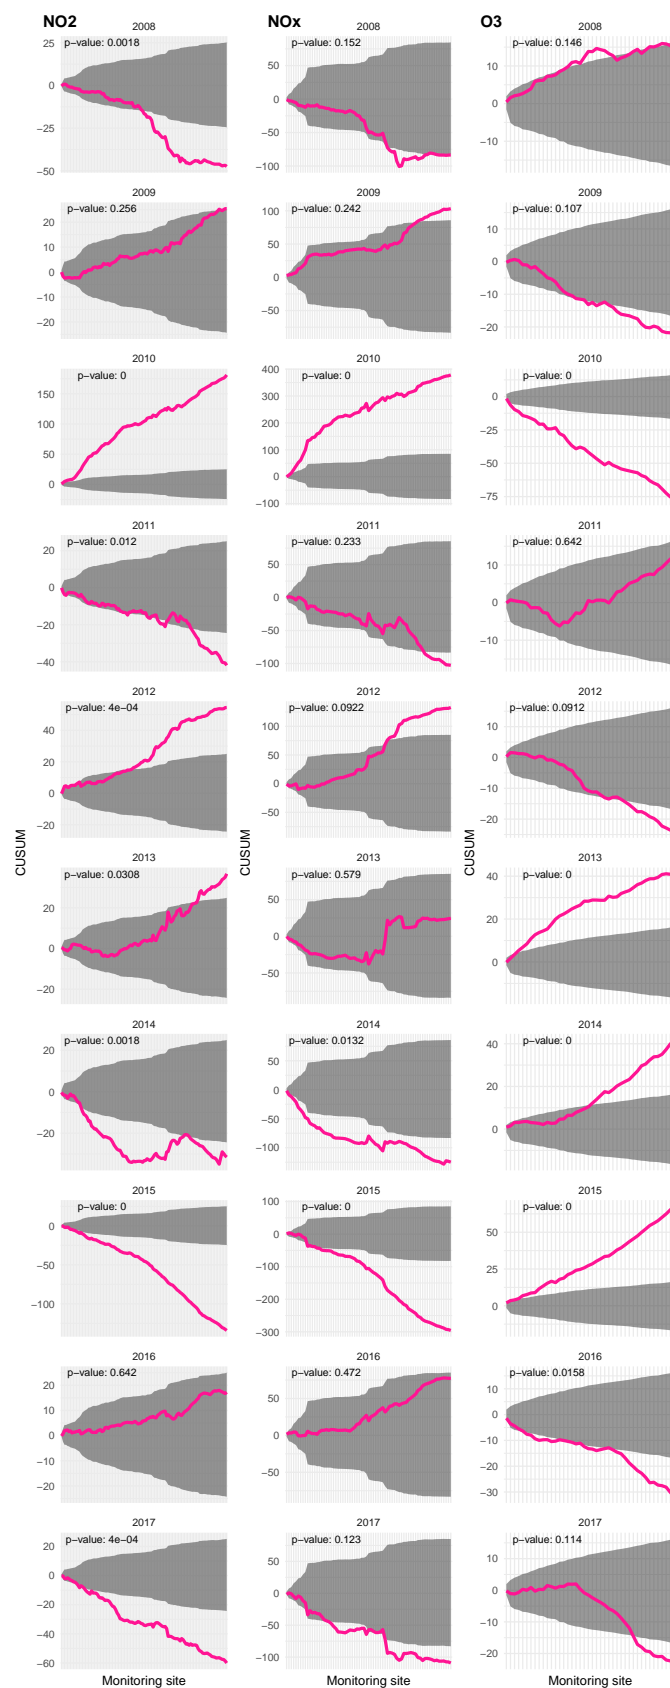
(a) Sites ordered by mean concentration

CHAPTER 4. QUANTIFYING THE EFFECT OF INTER-ANNUAL METEOROLOGICAL VARIATION ON POLLUTANT CONCENTRATIONS



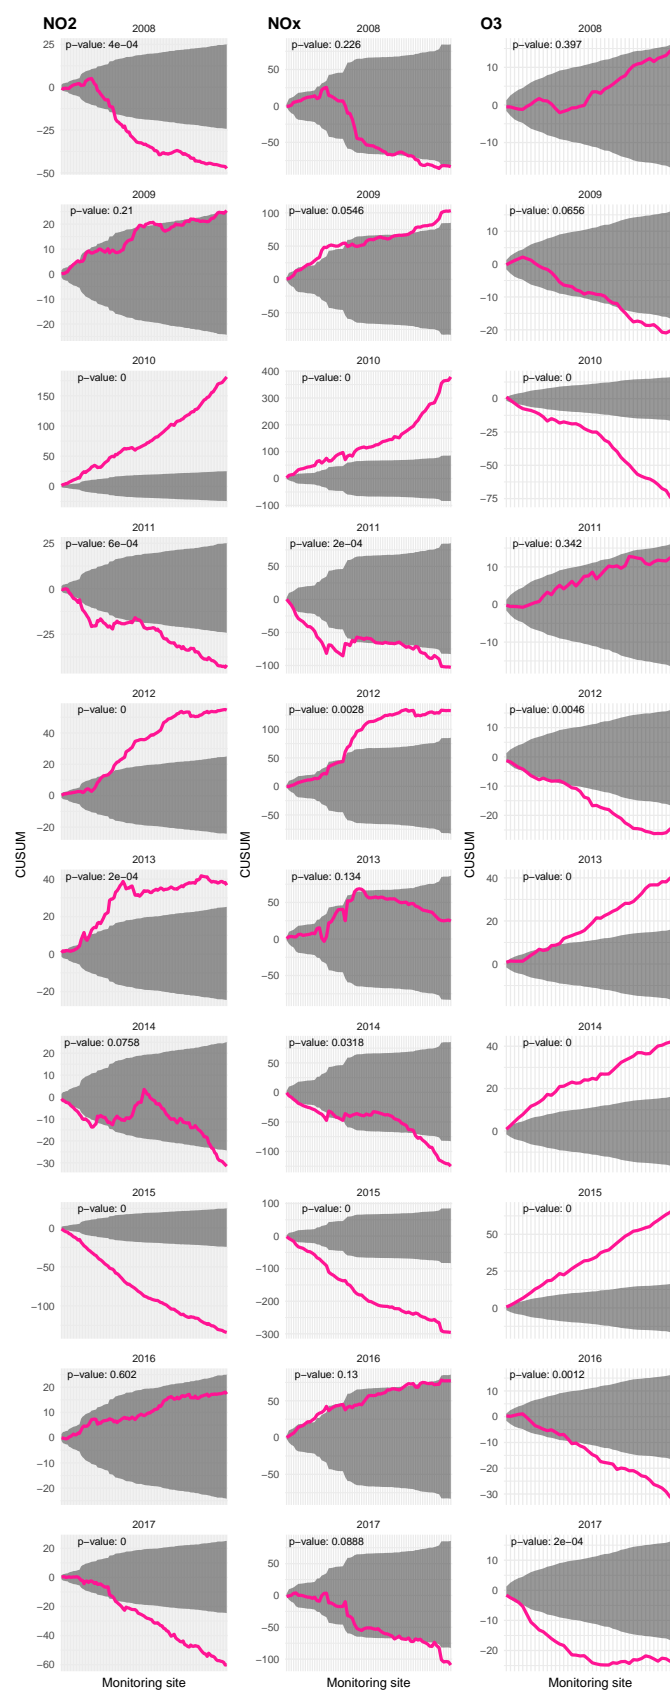
(b) Sites ordered by NO₂/NO_x ratio

CHAPTER 4. QUANTIFYING THE EFFECT OF INTER-ANNUAL METEOROLOGICAL VARIATION ON POLLUTANT CONCENTRATIONS



(c) Sites ordered by latitude (north to south from left to right)

CHAPTER 4. QUANTIFYING THE EFFECT OF INTER-ANNUAL METEOROLOGICAL VARIATION ON POLLUTANT CONCENTRATIONS



(d) Sites ordered by longitude (east to west from left to right)

Figure 4.12: Cumulative sum (CUSUM) plots of the difference between the mean concentration in a given meteorological year and the average concentration over all meteorological years for NO₂, NO_x and O₃ (columns 1-3 of each grid). The sites are ordered by (4.12a) mean pollutant concentration, (4.12b) NO₂/NO_x ratio, (4.12c) latitude (north to south from left to right), and (4.12d) longitude (east to west from left to right). Note: In the case of the O₃ CUSUM plots ordered by NO₂/NO_x ratio, only sites also measuring NO₂ and NO_x are used, and they are ordered by the NO₂/NO_x ratio derived from their measurements.

With the exception of the years 2017 and 2013, in all years the meteorology produces opposite effects on the concentrations of NO₂ and NO_x to that of O₃. For example, a year in which meteorology resulted in elevated NO_x and NO₂ concentrations generally also resulted in decreased O₃ concentrations, and vice versa. In 2017, however, the meteorology caused decreased concentrations of all pollutants, while in 2013 it drove elevated concentrations of all pollutants. One possible explanation for this is that meteorological conditions that are responsible for most NO₂ and NO_x pollution episodes, i.e. cold, dry, stable (low dispersion) winter weather would likely result in slower than normal ozone production. Ozone production is positively correlated with solar radiation and temperature, so in years with particularly low temperatures, ozone production, and therefore ambient concentration is observed to be lower.

For the most part, the effect of meteorology on pollutant concentrations was relatively consistent across monitoring sites with varying levels of pollution and contribution from traffic sources. The exceptions were a handful of sites with the very highest concentrations and the very lowest NO₂/NO_x ratios (i.e. the most polluted traffic sites — these included sites at London Marylebone Road, Knightsbridge, Chelsea, Earls Court Road, City of London (Walbrook Wharf), and Ealing (Hanger Lane Gyratory)), which sometimes displayed effects that differed from the consensus. For example, in 2011, 2014 and 2017, while the majority of sites experienced decreased NO₂ and NO_x concentrations due to meteorology, these sites were affected less. In 2016, while the majority of sites experienced increased NO₂ concentrations, these sites showed decreased NO₂ concentrations due to meteorology. As previously discussed, this is likely due to the relatively weaker influence of meteorology compared to emissions at these highly polluted roadside locations, as well as hyper-localised meteorology and street canyon effects.

For O₃, in 2008, 2010 and 2012, sites with low concentrations were less affected by meteorology (shallower slope). In contrast, in 2011 and 2017, sites with low concentrations were *more* affected by meteorology. For all other years, the slopes

were almost linear, indicating a similar influence of meteorology on all sites, regardless of their level of pollution.

The mean concentrations of NO_x and NO_2 at a site are (negatively) correlated with the NO_2/NO_x ratio (because sites with low NO_2/NO_x ratios tend to be characterised by strong traffic sources, which tend to result in high emissions and therefore high ambient concentrations of NO_x and NO_2), and therefore the results of the CUSUM analysis ordered by NO_2/NO_x ratio are similar to those of the analysis ordered by mean concentration for NO_2 and NO_x .

For O_3 , several of the CUSUM plots (2011, 2012, 2013, 2014, 2015) exhibit a curved shape, with the slope of the plot generally (except for 2013) greater for sites with high NO_2/NO_x ratio than low NO_2/NO_x . This indicates that the effect of meteorology on O_3 concentration is greater for sites with less influence from traffic sources (more 'background' sites).

More variation in the effect of meteorology was observed when the sites were ordered geographically. As with previous orderings, the CUSUM plots for NO_2 and NO_x tend to mirror each other fairly closely, and will therefore generally be discussed together.

In 2008, the meteorology had little effect on sites in the north and south of the UK, however it caused strong decreases in the concentrations at sites in the middle of the UK. O_3 concentrations were elevated across most sites, with the exception of several sites in the centre of the UK, which showed decreased concentrations. NO_2 and NO_x concentrations were elevated in the east, and depressed in the centre and west. O_3 concentrations were elevated across most of the UK, with the exception of a few sites in the centre, which showed reduced O_3 .

In 2009, elevated NO_x and NO_2 concentrations were observed at all sites, but particularly pronounced effects were observed for sites in the north and south, with less effect on central sites. O_3 concentrations were consistently lowered, with little variation either with latitude or longitude.

2010 was marked by highly elevated concentrations of NO_x and NO_2 at all sites, although the effect seemed to be slightly stronger at sites in the north and the west. Stronger effects of meteorology on O_3 were seen at sites in the east and west, with reduced effect in the centre. The Met Office described 2010 as "a year that was colder, drier and sunnier than average in most areas, particularly in the west." (Met Office, 2019) More extreme meteorology in the west of the country resulted in a greater effect on ambient concentrations.

In 2011, the decrease in NO_2 and NO_x concentrations, as well as the increased

O₃ concentrations, were observed at all sites, but were stronger in the south and east. O₃ concentrations were elevated at most sites, except a few sites in the east. The Met Office observed that in 2011, there was “a pronounced north-west to south-east gradient in rainfall with higher rainfall in the north-west, while drier than average in central, eastern and southern England” (Met Office, 2019).

An s-shape was observed for the 2012 CUSUMs, with the strongest effects of meteorology observed in the centre of the UK (both latitudinally and longitudinally) for all pollutants.

In 2013, meteorology had little effect on NO₂ concentrations in northern sites, while NO_x concentrations were decreased. In contrast, central and southern sites exhibited elevated concentrations of NO₂. Central sites, as well as the most southern sites showed little effect on NO_x, while centre-southern sites showed that meteorology caused increased NO_x concentrations. O₃ concentrations were most increased in the north, with smaller effects observed in the south. Sites in the east showed the strongest elevating effects of meteorology on NO_x and NO₂. Sites in central and western UK showed lowered NO_x concentrations. Generally, however, the meteorology in 2013 had a limited effect on concentrations. This is corroborated by the Met Office, which describes the meteorology as “generally near average and unremarkable” (Met Office, 2019).

In 2014, most sites showed decreased NO₂ and NO_x concentrations, with the greatest decreases at northern and western sites (and far eastern sites), however, a few central sites showed increased NO₂ concentrations. O₃ concentrations were relatively average at northern sites, with elevated O₃ at central and southern/western sites.

In 2015, the effects on meteorology on all pollutant concentrations were relatively consistent. A slightly greater effect on NO₂ and NO_x concentration was observed for southern sites than northern.

In 2016, greater elevations of NO₂ and NO_x concentration, as well as greater depressions in O₃ concentration were observed for southern sites, with a levelling off observed for the very most southern sites. The Met Office records that rainfall was lower than average in the south of the UK in 2016, possibly indicating the presence of anticyclonic conditions that led to build up of primary pollutants (Met Office, 2019).

In 2017, northern and western sites showed greater effects of meteorology on NO_x and NO₂ concentrations, while O₃ concentrations were relatively average at northern and western sites, and depressed in central/southern and eastern locations. The Met Office notes that the “north-east England sunniest relative

to average . . . it was rather drier across central and northern Scotland and many central and southern parts of England, but somewhat wetter in west Wales and north-west England. Sunshine was above average in the north and east, but slightly below average in some western fringes” (Met Office, 2019). The high rainfall in the north-west could have contributed to the greater decreases in NO_x and NO_2 .

Figure 4.13 summarises the results of the CUSUM analysis (with sites ordered by mean NO_2 concentration) for all three pollutants. The slope of the CUSUM plot represents the average per-site deviation from the average concentration, due to the meteorology experienced in that year. A positive slope indicates that the meteorology experienced in that year resulted in a higher-than-average ambient concentration of the pollutant in that year, while a negative slope indicates the reverse. The years on the x-axis are ordered by increasing slope of the NO_2 CUSUMs for clarity.




Figure 4.13: Summary of the results of the CUSUM analysis for NO_2 , NO_x and O_3 - the Theil-Sen slope of the CUSUM plot (ordered by mean pollutant concentration), which represents the average by-site deviation of the concentration due to meteorology from the average concentration 2008-2017. The meteorological years are ordered by their effect on the ambient NO_2 concentration, from the year which caused the largest decrease in NO_2 concentration on the left, to the year whose meteorology resulted in the greatest elevations of NO_2 concentration on the right.

Figure 4.13 makes three points apparent. First, that the meteorology in 2015, 2011, 2014, 2017 and 2008 resulted in decreased concentrations of NO_2 and NO_x , while that in 2009, 2013, 2016, 2012 and 2010 resulted in elevated concentrations of these pollutants. The meteorology caused decreased concentrations of ozone in 2017, 2009, 2016, 2012 and 2010 and increased concentrations in 2015, 2011, 2014, 2008 and 2013.

Second, that with the exception of the years 2017 and 2013, in all years the meteorology produces opposite effects on the concentrations of NO_2 and NO_x to that of O_3 . For example, a year in which meteorology resulted in elevated NO_x and NO_2 concentrations generally also resulted in decreased O_3 concentrations, and vice versa. In 2017, however, the meteorology caused decreased concentrations of all pollutants, while in 2013 it drove elevated concentrations of all pollutants.

Third, that 2010 stands out as a year in which the meteorology had a particularly dramatic effect on ambient concentrations of air pollutants, resulting in an average increase of $1.8 \mu\text{g m}^{-3}$ in NO_2 , an average increase of $6.85 \mu\text{g m}^{-3}$ in NO_x , and an average decrease of $2.2 \mu\text{g m}^{-3}$ at each monitoring site.

4.3.3 The Relationship between Ambient Concentration and Annual Meteorology

The year-to-year variations in air pollutant concentrations can be interpreted in terms of the inter-annual variations in meteorology. For this purpose, the relationships between meteorology and pollutant concentration were investigated in more detail, in order to ascertain the specific meteorological patterns responsible for driving the observed variation in the annual average concentrations of the pollutants. Figure 4.14 shows the values of each of the meteorological variables used in model training, as a function of the Julian day, for each year analysed. These annual distributions are compared to the average daily values of the meteorological variables over all years in the analysis (i.e. the 'average' meteorological conditions), which are shown by the yellow lines in Figure 4.14. The information in Figure 4.14 was supplemented with information from the Met Office annual climate summaries (Met Office, 2019).

CHAPTER 4. QUANTIFYING THE EFFECT OF INTER-ANNUAL METEOROLOGICAL VARIATION ON POLLUTANT CONCENTRATIONS




Figure 4.14: A comparison of the average UK meteorological conditions (air temperature, air pressure, relative humidity, dew point, wind speed, wind direction and visibility) as a function of the Julian day in each year 2008–2017 (purple smooth lines) and the average UK meteorological conditions over the entire ten-year period (yellow smooth lines). All hourly data used as predictors during random forest model training were averaged to daily data, and the line represents a loess smooth through the daily averages. The shaded band represents the 95% confidence interval around the loess smooth.

2010 and 2012 exhibited significantly elevated concentrations of NO_2 and NO_x due to the meteorology in those years. 2010 was described by the Met Office as exhibiting “Prolonged periods with blocked weather patterns, and an associated absence of westerly airstreams, resulted in a year that was colder, drier and sunnier than average in most areas, particularly in the west.” 2012 was described as having an “exceptionally wet period for most of the country from April lasting through much of the summer.” The Met Office also mentions that “It is worth noting that only 2 years (2010 and 2012) of the last 16 have had annual temperatures below this average.” (Met Office, 2019) From this, as well as from Figure 4.14, it can be seen that both 2010 and 2012 were characterised by cold

temperatures and low dispersion conditions in the winter. This resulted in higher concentrations of primary pollutants and pollution episodes, due to the lack of mixing. The low temperatures were also likely responsible for slower production of secondary ozone, resulting in the observed decreased ozone concentration due to meteorology in these years.

In contrast, the years 2011, 2014 and 2015 showed lowered NO_x and NO_2 concentrations, and elevated O_3 concentrations due to meteorology. These years were characterised by warmer than average temperatures, high rainfall and high dispersion conditions (wet and windy) in the winter. The Met Office records that 2011 was “the second warmest year in the series from 1910”, in 2014 “all months except August were warmer than average, and this was the warmest year on record for the UK. It was also wetter than average for many locations, . . . the winter storms of January and February, which brought damaging winds, with inland and coastal flooding.” 2015 was described as “The summer was rather cool and wet, but early autumn provided fine, sunny weather as compensation. However from late autumn a succession of Atlantic storms brought exceptional rainfall to the north and west, causing widespread severe flooding to many towns and cities . . . The UK mean temperature was 9.2 °C, 0.4 °C above the 1981–2010 long term average . . . The UK rainfall total was 1272 mm, 110% of the 1981–2010 average and seventh-wettest in the UK series” (Met Office, 2019). The wet and windy weather indicates that low pressure cyclonic conditions dominated in these years, leading to increased dispersion of atmospheric pollution.

The Met Office described 2013 as “A late winter and exceptionally cold spring, with unseasonably late snowfalls, lead into a warm and sunny summer. October and December saw Atlantic storms that brought rain and at times very high winds, causing widespread disruption.” (Met Office, 2019) This is corroborated by Figure 4.14, which shows lower-than-average temperatures in the spring and higher than average temperatures in summer, as well as high wind speeds in spring and winter (November - December). This mix of high and low dispersion conditions throughout the year likely contributed to the mixed effects on air quality observed in the CUSUM plots for 2013 (p-values were often >0.05 , especially for NO_x , indicating no significant effect of meteorology on concentrations).

The meteorology in 2017 resulted in lowered concentrations of NO_x , NO_2 and O_3 . This year was marked by warmer than average temperatures, and average rainfall. It was sunnier than average, with a weather event characterised by particularly high temperatures and sunshine in June. The Met Office described it as “The year as a whole was rather warmer than average for the UK . . . 2017 was

a slightly sunnier than average year for the UK as a whole . . . Notable extreme weather events during the year included Storm Doris in February and flash flooding in Coverack, Cornwall in July; autumn and early winter saw occasional notable storm systems . . . The hot spell in June saw the highest temperatures in that month for over 40 years, and, unusually, brought temperatures above 30 °C somewhere in the UK five days in a row.” (Met Office, 2019) The weather events could have contributed to the lowered concentrations of primary pollutants, as they may have resulted in increased dispersion. The warm temperatures may also have prevented winter blocking episodes from causing pollution episodes. However, since warm, sunny conditions increase the rate of secondary ozone production, it is unexpected that the year would exhibit lowered ozone concentrations.

4.4 Conclusions and Future Directions

Inter-annual variations in meteorology were found to have a considerable effect on ambient concentrations of NO₂, NO_x and O₃. On average, the range of the annual average concentrations of these species due to meteorological variation was 2.9 µg m⁻³ (8.2%) for NO₂, 9.9 µg m⁻³ (12.6%) for NO_x and 3.3 µg m⁻³ (7.5%) for O₃.

This uncertainty caused by meteorological variation has implications for the evaluation of compliance with EU limit values. The EU limit value for the annual average NO₂ concentration is 40 µg m⁻³. Of the 173 monitoring sites included in the analysis, an average of 54.8 sites exceeded this value in any given year. The average number of ‘marginal’ sites for which the EU limit value for annual mean NO₂ concentration lay within the range resulting from meteorological variation was 12.1, or 22% of the number of sites exceeding the limit value. In total, 44 monitoring sites were ‘marginal’ in at least one year over the period of analysis, or 25% of the total number of sites. In other words, in any given year, the compliance of around 22% of the monitoring sites was dependent on the meteorology in that year. For example, a marginal site may be ‘compliant’ with the EU limits in a year in which meteorology acts to reduce the concentration of NO₂, while it may exceed the limit value in a year in which meteorology has the reverse effect. Whether or not the site exceeds the limit value is, under these circumstances, dependent on variations in the local weather, which is clearly counter to the aim of compliance monitoring.

The heatmaps and CUSUM analysis showed that the meteorology in 2015, 2011, 2014, 2017 and 2008 resulted in decreased concentrations of NO₂ and NO_x

for the majority of UK monitoring sites, while that in 2009, 2013, 2016, 2012 and 2010 resulted in elevated concentrations of these pollutants. In most years, meteorology produced opposite effects on the concentration of O_3 to that of NO_2 and NO_x , with decreased concentrations of ozone in 2017, 2009, 2016, 2012 and 2010 and increased concentrations in 2015, 2011, 2014, 2008 and 2013.

2010 stands out as a year in which the meteorology had a particularly dramatic effect on ambient concentrations of air pollutants, resulting in an average increase of $1.8 \mu\text{g m}^{-3}$ in NO_2 , an average increase of $6.9 \mu\text{g m}^{-3}$ in NO_x , and an average decrease of $2.2 \mu\text{g m}^{-3}$ at each monitoring site.

Low dispersion, anticyclonic conditions characterised by cold temperatures and low dispersion conditions in the winter were responsible for driving elevated concentrations of NO_x and NO_2 , and lowered concentrations of O_3 , as a result of a lack of mixing of primary pollutants (with long periods of stability sometimes resulting in pollution episodes), and slower production of secondary ozone.

High dispersion, cyclonic conditions characterised by warmer than average temperatures, high rainfall and high dispersion conditions (wet and windy) in the winter, on the other hand, resulted in decreased concentrations of NO_2 and NO_x , and usually elevated concentrations of ozone.

In general, different monitoring sites exhibited similar dependencies on meteorology. However, some of the most polluted roadside monitoring sites exhibited slightly different effects from meteorology, possibly due to the weaker dependence of the ambient concentration on meteorology, or the inability of the meteorological measurements to represent the hyper-local conditions often present at such sites. In addition, the CUSUM plots were able to reveal geographical differences in the effects of meteorology on ambient concentrations, resulting from differences in local meteorology across the UK. The detailed interpretation of these spatial differences may be the target of future work.

4.4.1 Implications for Compliance Monitoring and Air Quality Modelling

The results of the analysis, and the considerable influence that meteorology was found to have upon the annual average concentrations of NO_2 , NO_x and O_3 , suggest that the annual mean is not a robust means of assessing compliance. A more robust metric would be able to provide information about the level of pollution at the monitoring site without being sensitive to inter-annual variations in local meteorology. One possible method could involve the use of the random

forest methodology to 'de-weather' the ambient data. The annual average de-weathered concentration, or the three-year rolling mean, could then be assessed against a limit value. Alternatively, the trend in the ambient concentration, rather than the annual mean, could be assessed.

However the use of these methods would necessitate greater computational resource and specialist knowledge than a simple average. Furthermore, a greater amount of data would be required – ideally the model would be trained using at least three years of data.

Meteorology appears to have exerted the most extreme effects on ambient concentrations of NO_2 , NO_x and O_3 in 2010, 2015, 2011, 2014, 2012, 2016 and 2017. These years are therefore likely to lead to biased results if used as baseline years in air quality modelling. For example, the use of 2015 as the baseline year would yield an overly-optimistic view of air quality. A more appropriate choice for a baseline year would be one of the years in which meteorology exerted the smallest influence on concentrations, such as 2008, 2009, or 2013.

4.4.2 Future Directions

In addition to analysis of spatial differences in the effects of meteorology on air quality, and the identification of more robust methods for assessing compliance, future work may also focus on quantifying the range imposed by inter-annual variations in meteorology on the concentrations of other important regulated species, such as PM_{10} and $\text{PM}_{2.5}$. Like NO_2 , annual mean concentrations are used to measure compliance with limit values for both of these pollutants, and therefore may also be sensitive to these effects.

Future work may also focus on the development and application of the CUSUM methodology. For example, different monitoring site orderings of the CUSUM plots, such as altitude for ozone, may also reveal additional information. The method, combined with a geographical ordering, also has the potential to reveal patterns in the spatial distributions of pollutant concentration, and changes in the concentration, which may otherwise not be easy to visualise.

Bibliography

AQEG, 2004. Nitrogen dioxide in the United Kingdom. Tech. rep., Air Quality Expert Group.

URL <https://uk-air.defra.gov.uk/library/aqeg/publications>

Carslaw, D., 2017. worldmet: Import Surface Meteorological Data from NOAA Integrated Surface Database (ISD).

URL <https://github.com/davidcarslaw/worldmet>

Carslaw, D., 2020. Ricardo: an analysis of Covid-19 lockdown on UK local air pollution.

URL <https://airqualitynews.com/2020/03/31/ricardo-an-analysis-of-covid-19-lockdown-on-uk-local-air-pollution/>

Defra, 2019. UK and EU Air Quality Limits.

URL <https://uk-air.defra.gov.uk/air-pollution/uk-eu-limits>

Grange, S. K., Carslaw, D. C., Lewis, A. C., Boleti, E., Hueglin, C., 2018. Random forest meteorological normalisation models for Swiss PM₁₀ trend analysis. Atmospheric Chemistry and Physics Discussions, 1–28.

URL <https://doi.org/10.5194/acp-18-6223-2018>

Jiang, N., Dirks, K. N., Luo, K., 2014. Effects of local, synoptic and large-scale climate conditions on daily nitrogen dioxide concentrations in Auckland, New Zealand. International Journal of Climatology 34 (6), 1883–1897.

URL <http://doi.wiley.com/10.1002/joc.3808>

Manly, B. F., Mackenzie, D. I., 2003. CUSUM environmental monitoring in time and space. Environmental and Ecological Statistics 10 (2), 231–247.

URL <http://link.springer.com/10.1023/A:1023682426285>

Manly, B. F. J., Mackenzie, D. I., 2000. A cumulative sum type of method for environmental monitoring. Environmetrics 11 (2), 151–166.

BIBLIOGRAPHY

URL <https://www.researchgate.net/publication/235961584-A-cumulative-sum-type-of-method-for-environmental-m>

Met Office, 2019. Climate summaries.

URL <https://www.metoffice.gov.uk/climate/uk/summaries>

NOAA, 2016. Integrated Surface Database.

Wayne, R. P., 2000. Chemistry of Atmospheres, 3rd Edition. Oxford University Press.

5. Conclusions & Future Directions

Air pollution is arguably the world's largest environmental health threat. In 2016, 91% of the global population were living in places where air pollution exceeded the World Health Organisation's guidelines (WHO, 2020b). In Europe alone, exposure to PM decreases the life expectancy of each person, on average, by almost 1 year (WHO, 2020a).

Clearly, it is vitally important to implement technologies, policies and behavioural changes that reduce the emissions of air pollutants. But emissions are only part of the story. The complexity of atmospheric chemistry, the unpredictable influence of meteorology, and the potential for error in accounting for emissions (as emphasised by the recent Volkswagen emissions scandal (Lewis et al., 2015)) means that it is equally important to carefully monitor and analyse ambient air quality and human exposure.

In this thesis, several statistical techniques have been developed and applied to routine ambient monitoring data collected in the UK to overcome some of the difficulties associated with the data set and to extract further insight. In particular, these techniques have focused on the analysis of data from *networks* of monitoring sites, in order to obtain information about large-scale influences on air quality such as the effects of changes in vehicle emissions.

The problems associated with conducting trend analysis on a large scale using networks of multiple monitoring sites with time series of differing lengths were examined in Chapter 2, and the rolling change method was developed in response to these issues. The efficacy of the method was established using simulated data, and by comparing its results with those obtained using standard methods and the results of other studies.

The rolling change method was used to analyse the long term trends in the roadside concentrations of NO_x , NO_2 and PM_{10} , as well as the NO_2/NO_x ratio, between 2000 and 2017 across Scotland and the UK. Similarly shaped trends were observed at each scale. NO_x concentrations decreased monotonically throughout

the period of study, and after an initial period of stability or slight increase, NO₂ concentrations also declined monotonically. These changes are most likely the result of the introduction of vehicle exhaust technologies aimed at reducing emissions of these species, such as three-way catalysts on petrol vehicles, and Lean NO_x Traps (LNT) and Selective Catalytic Reduction (SCR) on diesel vehicles. The initial increase in NO₂ concentration was probably the result of increased vehicular emissions of NO₂, due to the introduction of Diesel Oxidation Catalysts (DOC) and Diesel Particulate Filters (DPF) to diesel vehicles, which deliberately oxidise NO to NO₂ for use in the oxidation of other pollutants, such as CO, hydrocarbons and particulate matter. The increase in the proportion of diesel vehicles in the vehicle fleet during this time could also have been a contributing factor.

The NO₂/NO_x ratio increased until around 2010, before decreasing. This initial increase can be explained by slower decline in the NO₂ concentration than in the NO_x concentration, resulting from the introduction of DOC and DPF, as previously described. The decrease in NO₂/NO_x ratio observed since around 2010 is likely the result of a combination of factors, including the modification of vehicle emission after-treatment systems to no longer over-produce NO₂, and a decrease in the NO₂/NO_x ratio of diesel vehicle emissions as the vehicle mileage increases, an effect that has recently been observed in vehicle emission remote sensing measurements (Carslaw et al., 2019).

The PM₁₀ concentration was initially relatively constant, before declining. This decline can be linked to the fitting of many Euro 4 vehicles, and all post-Euro 5 vehicles, with diesel particulate filters (DPF).

While the shapes of these trends were very consistent across different scales and locations, the timing of the turning points occasionally differed. The turning points in Scotland generally occurred later than those across the whole UK, indicating that the decline in the vehicular emissions of these species began later than average. This could be due to differences in the vehicle fleet in Scotland compared with London and the UK.

The development of the rolling change method is timely given that low cost sensor networks are increasingly being used to monitor air quality. These sensors tend to be less reliable than traditional monitoring sites, as well as being more portable, and therefore such networks are even more vulnerable to the problems associated with site flux. As such, the rolling change method could be valuable in analysis of trends using data from these networks. Moreover, the issue of aggregating data from time series of differing lengths is not exclusive to the

analysis of air quality data. The method could find application in any discipline dealing with multiple time series of differing lengths.

To facilitate the use of the rolling change method, an open source R package named `aqtrends` was developed. The package contains a function that outputs the rolling trend for a given data set input. More details about `aqtrends` is given in Appendix III.

In Chapter 3, random forest models were used to model ambient air quality as a function of background concentrations, local meteorology and temporal variables was investigated. The models were used to remove the influence of confounding factors enabling a clearer view of the trends in air pollutant concentration. The de-weathered trends in the concentrations of NO_x , NO_2 and PM_{10} at three London monitoring sites (Marylebone Road, Camden Kerbside and Cromwell Road) between 2000 and 2017 were calculated.

At all three sites, PM_{10} concentrations decreased from 2007, coinciding with the penetration of Euro 4 vehicles (which were the first to be fitted with DPF) into the fleet. The eventual levelling off of the PM_{10} concentration observed at one site was possibly due to saturation of the fleet with Euro 4 and above vehicles. However, any interpretation must consider the difficulty of analysing PM_{10} concentrations, even when meteorology is accounted for, due to the numerous other sources that may affect the concentration. For this reason, it may be advantageous to focus on specific PM components, such as black carbon or particle number, in future work.

The trends in NO_x concentration differed between the sites, possibly due to local influences. NO_2 concentrations were found to generally decrease, in consistency with decreases in vehicle emissions from the introduction of Euro 4 and 5 vehicles.

An investigation of the effects of the implementation of the London LEZ on the ambient roadside concentration of PM_{10} at a number of London monitoring sites was also conducted, using the random forest models to model air pollutant concentrations as a function of local conditions, as well as an additional indicator variable representing the stage of the intervention (Phase 0 indicating the period prior to the implementation of the LEZ, through to Phase 3, the final and most restrictive phase of the LEZ). De-weathered trends were calculated for each value of the indicator variable (i.e. for each stage of the intervention), in order to generate counter-factual scenarios representing the trends that would have been observed had a single phase of the intervention been in force throughout the duration of the period. The resulting trends indicated that, while all sites showed

a decrease in PM₁₀ concentration between the 'no LEZ' and 'final phase LEZ' scenarios, the results of the consecutive phases were more mixed, with some sites showing increases in PM₁₀ concentration for Phases 1 and 2. The fact that no clear evidence of any impact on air quality due to the LEZ despite the sensitivity of the method suggested that the raw data used in this analysis is insufficient to detect the changes. Inclusion of further data in future analyses, as well as considering the increment about background concentration rather than raw concentration, may produce more definitive insights in future work. Additionally, the earlier comments regarding the advantages of using specific PM components also apply here.

Similar methodologies could be used in the future to analyse other interventions and drivers of changes in air quality. For example, much has been made recently of the improvements in air quality in urban areas that accompanied the Covid-19 lockdown (He et al., 2020). The methods described in Chapter 3 would be well-suited to assessment and quantification of the 'true' impact of the lockdown, independent of the effects of other confounding factors.

In Chapter 4, the effects of inter-annual meteorological variation on the annual mean concentrations of NO_x, NO₂ and O₃ were estimated and quantified using a variation of the de-weathering method described in Chapter 3. Inter-annual variations in meteorology were found to have a considerable effect on ambient concentrations of NO₂, NO_x and O₃. On average, the range of the annual average concentration due to meteorological variation was 2.9 µg m⁻³ (8.2%) for NO₂, 9.9 µg m⁻³ (12.6%) for NO_x and 3.3 µg m⁻³ (7.5%) for O₃.

These results have implications for compliance monitoring. Compliance with EU limit values for some pollutants, such as NO₂, is evaluated using the annual mean metric. For example, the EU limit value for the annual average NO₂ concentration is 40 µg m⁻³. Of the 173 monitoring sites included in the analysis, an average of 54.8 sites exceeded this value in any given year. The average number of 'marginal' sites for which the EU limit value for annual mean NO₂ concentration lay within the uncertainty resulting from meteorological variation was 12.1, or 23% of the number of sites exceeding the limit value. In total, 44 monitoring sites were marginal in at least one year over the period of analysis, or 25% of the total number of sites. In other words, in any given year, the compliance of around 23% of the monitoring sites was dependent on the meteorology in that year. These findings suggest that the annual mean metric is not a robust way for measuring compliance, as whether or not a site is found to be compliant in any given year is

often dependent on the weather experienced during that year.

CUSUM plots and heatmaps proved to be an effective way of visualising differences in the effects of meteorology on air pollutant concentration across a monitoring network of sites. These tools enabled both the visualisation of differences in the effects on concentration *between* meteorological years, and differences in the effects on concentration *within* years across the network based on the ordering variable, e.g. latitude or longitude.

A useful feature of the CUSUM analysis is the ability to order the monitoring sites by different variables, such as the mean concentration, NO_2/NO_x ratio, latitude and longitude, revealing subtleties in the relationships between meteorology and air quality across the individual sites. In general, most monitoring sites exhibited similar dependencies on meteorology, however, some of the most polluted roadside monitoring sites exhibited slightly different effects from meteorology, possibly due to the weaker dependence of the ambient concentration on meteorology, or the inability of the meteorological measurements to represent the hyper-local conditions often present at such sites. In addition, the CUSUM plots were able to reveal geographical differences in the effects of meteorology on ambient concentrations, resulting from differences in local meteorology across the UK. These tools may be useful in conducting or interpreting future analyses of monitoring networks.

5.1 Future Directions

The work described in this thesis could be developed in several ways.

In the future, the rolling change method may be applied to extract insight regarding large scale trends in air pollution at many different scales, and in different locations. The problem of site flux is common in air quality monitoring networks worldwide, and the rolling change method could enable long term trend analysis in sparse monitoring networks without many, or any, long running sites. The *aqtrends* R package was created to facilitate the use of the method in future work (Lang, 2018).

Additionally, future work might extend the rolling change method to include estimates of uncertainty in the trend.

In Chapter 4, the range imposed by the inter-annual meteorological variation on the concentrations of air pollutants was only calculated for NO_x , NO_2 and O_3 . Future work may extend the analysis by also considering other important regulated species, such as PM_{10} and $\text{PM}_{2.5}$. Since annual mean concentrations are

also used to measure compliance with limit values for both of these pollutants, it would be important to quantify the sensitivity of their mean values to inter-annual variations in meteorology.

Further development and application of the CUSUM methodology might also be a promising avenue for future studies. For instance, when ordering monitoring sites by geographical location, it could be a valuable visualisation tool for revealing patterns and changes in the spatial distributions of pollutant concentrations. Furthermore, the method offers great flexibility in the choice of variables by which to order the monitoring sites. In the work conducted in Chapter 4, variations in concentration across the network of monitoring sites were only considered as a function of mean concentration at the site, average NO₂/NO_x ratio (relative contribution from the traffic source), latitude and longitude. However, the use of different variables, such as altitude, by which to order the sites, could reveal more insights.

Given the conclusions of the work in Chapter 4, that inter-annual variations in meteorology exert a considerable influence on compliance or exceedance with annual mean concentration limits at many UK monitoring sites, more thought should be given to alternative ways of measuring compliance. For example, using a rolling mean to measure NO₂ concentration. Future work could attempt to ascertain whether different metrics reduce the influence of inter-annual meteorological variation.

Bibliography

Carslaw, D. C., Farren, N. J., Vaughan, A. R., Drysdale, W. S., Young, S., Lee, J. D., 2019. The diminishing importance of nitrogen dioxide emissions from road vehicle exhaust. *Atmospheric Environment: X* 1, 100002.

URL <https://doi.org/10.1016/j.aeaoa.2018.100002>

He, G., Pan, Y., Tanaka, T., 2020. The short-term impacts of COVID-19 lockdown on urban air pollution in China. *Nature Sustainability* 3 (7).

URL <https://www.nature.com/articles/s41893-020-0581-y>

Lang, P., 2018. aqtrends.

BIBLIOGRAPHY

URL <https://github.com/pollylang/aqtrends>

Lewis, A. C., Carslaw, D. C., Kelly, F. J., 2015. Diesel pollution long under-reported. Nature 526, 195.

URL <http://dx.doi.org/10.1038/{}0A526195c>

WHO, 2020a. Air quality: Data and statistics.

URL <https://www.euro.who.int/en/health-topics/environment-and-health/air-quality/data-and-statistics>

WHO, 2020b. Ambient (outdoor) air pollution.

URL [https://www.who.int/news-room/fact-sheets/detail/ambient-\(outdoor\)-air-quality-and-health](https://www.who.int/news-room/fact-sheets/detail/ambient-(outdoor)-air-quality-and-health)

A. Appendix I: Algorithm for Chapter 2

The algorithm for the rolling change method described in Section 2.2.2 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_1 as the average of the annual average concentrations of all monitoring sites in the first year, y_1 .
3. Identify the moving window, i , as the time period $x_1, \dots, x_{1+(n-1)}$.
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 2.1.
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 2.2.
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 2.3.

B. Appendix II: Supplementary material for Chapter 2

Table B.1 contains information about the London roadside monitoring sites measuring NO_x and NO_2 between 2000 and 2017 that were used in the trend analysis described in Section 3.2.

APPENDIX B. APPENDIX II: SUPPLEMENTARY MATERIAL FOR CHAPTER 2

Table B.1: Metadata for London roadside monitoring sites used in the analysis. The mean and standard deviation of the hourly NO_x and NO₂ concentration data measured by each site during the period of 2000 to 2017 is also shown. Monitoring sites were selected as those within the bounding box of coordinates 51.25°N, 51.71°N, -0.54°E, 0.28°E. The data was sourced from the AURN, LAQN and AQE networks.

site code	site name	latitude	longitude	altitude	network id	[NO _x] (µg m ⁻³)		[NO ₂] (µg m ⁻³)	
						mean	SD	mean	SD
A3	London A3 Roadside	51.37	-0.29	32.00	aun	173.00	144.00	60.70	30.90
A30	Kingston - Kingston Bypass A3	51.37	-0.29	27.89	laqn	173.00	144.00	60.60	30.80
BG3	Barking and Dagenham - North Street	51.54	0.07	6.29	laqn	102.00	106.00	49.70	29.00
BN1	Tally Ho	51.61	-0.18		aqe	167.00	113.00		
BN2	London Barnet Chalgrove School	51.59	-0.21		aqe	62.70	87.90	33.50	22.70
BRN	Brentford Roadside	51.49	-0.31	15.00	aun	151.00	133.00	53.80	29.60
BT2	Brent - Ikea Car Park	51.55	-0.26	31.69	laqn	296.00	205.00	67.30	28.80
BT3	Brent - Harlesden	51.54	-0.25	48.64	laqn	118.00	99.10	52.50	27.10
BX7	Bexley - Thames Road North	51.46	0.20	8.25	laqn	101.00	106.00	42.40	23.30
BX8	Bexley - Thames Road South	51.46	0.19	9.02	laqn	84.50	94.70	39.10	25.40
BY2	London Bromley	51.41	0.02	50.00	aun	105.00	89.10	49.40	24.80
BY7	Bromley - Harwood Avenue	51.41	0.02	65.19	laqn	98.60	84.10	48.60	24.80
CA1	Camden Kerbside	51.54	-0.18	50.00	aun	182.00	148.00	69.80	37.20
CD1	Camden - Swiss Cottage	51.54	-0.18	57.25	laqn	182.00	148.00	69.50	37.20
CD2	Camden API	51.54	-0.18	57.13	laqn	96.70	83.00	54.80	38.60
CD3	Camden Shaftesbury Avenue	51.51	-0.13		aqe	169.00	109.00	76.60	29.40
CD9	Camden - Euston Road	51.53	-0.13	22.55	laqn	325.00	210.00	100.00	46.30
CR2	Croydon - Purley Way	51.36	-0.12	52.85	laqn	141.00	109.00	45.40	24.10
CR4	Croydon - George Street	51.37	-0.10	58.38	laqn	108.00	87.00	51.50	26.80
CR5	Croydon - Norbury	51.41	-0.12	36.46	laqn	184.00	157.00	64.70	39.10
CR7	Croydon - Purley Way A23	51.36	-0.12	54.95	laqn	84.20	72.90	33.40	19.40
CR9	Croydon - Park Lane	51.37	-0.10	58.60	laqn	137.00	114.00	49.80	29.40
CRD2	London Cromwell Road 2	51.50	-0.18	20.00	aun	181.00	111.00	75.10	28.40
CT6	City of London - Walbrook Wharf	51.51	-0.09	21.08	laqn	343.00	281.00	112.00	49.70
CY1	Crystal Palace - Crystal Palace Parade	51.42	-0.08	109.46	laqn	122.00	99.40	48.40	25.00
EA2	Ealing - Acton Town Hall	51.51	-0.27	17.71	laqn	143.00	131.00	57.00	30.70
EA6	Ealing - Hanger Lane Gyrotory	51.53	-0.29	37.51	laqn	301.00	235.00	86.10	44.30
EI1	Ealing - Western Avenue	51.52	-0.27	30.09	laqn	159.00	132.00	62.80	31.00
EI2	Ealing - Southall Railway	51.51	-0.38	31.49	laqn	84.30	83.90	34.80	20.10
EL2	Elmbridge - Esher High Street	51.37	-0.36	31.78	laqn	114.00	97.20	47.70	25.70
EL3	Elmbridge - Hampton Court Parade	51.40	-0.34	8.92	laqn	155.00	124.00	55.60	27.40
EN2	Enfield - Church Street	51.65	-0.08	29.51	laqn	85.60	87.20	42.60	23.10
EN4	Enfield - Derby Road	51.61	-0.05	10.97	laqn	104.00	101.00	46.20	21.50
EWE2	Ewell High Street	51.35	-0.25		aqe			44.80	26.90
FA1	Team London Bridge - Tooley Street FASQ	51.50	-0.08	4.17	laqn			62.60	33.70
GB6	Greenwich and Bexley - Falconwood	51.46	0.09	64.45	laqn	120.00	123.00	46.00	31.50
GN0	Greenwich - A206 Burrage Grove	51.49	0.07	8.76	laqn	89.80	96.60	45.80	27.60
GN3	Greenwich - Plumstead High Street	51.49	0.10	12.20	laqn	80.90	76.80	39.70	25.20
GN4	Greenwich - Fiveways Sidcup Rd A20	51.43	0.06	55.69	laqn	128.00	114.00	48.90	29.40
GR5	Greenwich - Trafalgar Road	51.48	-0.00	7.01	laqn	96.60	82.10	47.30	25.10
GR7	Greenwich - Blackheath	51.47	-0.01	21.73	laqn	113.00	92.50	46.40	24.60
GR8	Greenwich - Woolwich Flyover	51.49	0.02	3.82	laqn	221.00	176.00	70.90	36.70
GR9	Greenwich - Westthorne Avenue	51.46	0.04	30.26	laqn	102.00	114.00	43.40	26.80
HB102	Broxbourne	51.68	-0.03		aqe			46.30	26.10
HB126	Watford Roadside	51.66	-0.40		aqe			40.40	22.20
HF1	Broadway	51.49	-0.22		aqe	221.00	156.00	76.80	37.70
HF4	Shepherd's Bush	51.50	-0.22		aqe			80.00	36.30
HFY	Hammersmith and Fulham - Talgarth Road B	51.50	-0.08	4.17	laqn			61.20	35.00
HG1	Haringey Roadside	51.60	-0.07	10.00	aun	96.20	90.90	44.00	22.10
HG3	Haringey - Bounds Green	51.61	-0.12	47.90	laqn	125.00	131.00	51.50	26.60
HI1	Hillingdon 1 - South Ruislip	51.55	-0.40		aqe	121.00	117.00	46.00	26.40
HIL5	London Hillingdon Hayes	51.50	-0.41	30.00	aqe			47.30	26.10
HK6	Hackney - Old Street	51.53	-0.08	16.98	laqn	147.00	86.60	62.80	24.40
HR2	Harrow - Pinner Road	51.59	-0.36	52.40	laqn	111.00	107.00	46.10	27.40
HS1	Hounslow Roadside	51.49	-0.31	15.00	aun	144.00	137.00	54.50	25.70
HS2	Hounslow Cranford	51.48	-0.41		aqe	68.30	86.80	35.90	22.30
HS6	Hounslow Heston	51.48	-0.36		aqe	123.00	108.00		
HS8	Hounslow Gunnersbury	51.50	-0.28	2.00	aqe	172.00	136.00	57.20	31.40
HV1	Havering - Rainham	51.52	0.21	8.10	laqn	83.00	88.20	39.40	23.40
HV3	Havering - Romford	51.57	0.18	14.37	laqn	88.70	92.10	39.70	22.00
IM1	Camden - Holborn (Bee Midtown)	51.52	-0.12	27.28	laqn	235.00	178.00	83.90	39.70
IS2	Islington - Holloway Road	51.56	-0.12	28.94	laqn	163.00	109.00	62.60	27.20
KC2	Cromwell Road	51.50	-0.18		aqe	167.00	108.00	70.30	28.60
KC3	Knightsbridge	51.50	-0.16		aqe	214.00	185.00	85.20	50.00
KC4	Chelsea	51.49	-0.17		aqe	219.00	144.00	86.60	38.70
KC5	Earls Court Road	51.49	-0.19		aqe	277.00	187.00	96.70	50.40

APPENDIX B. APPENDIX II: SUPPLEMENTARY MATERIAL FOR CHAPTER 2

site code	site name	latitude	longitude	altitude	network id	[NO _x] (µg m ⁻³)		[NO ₂] (µg m ⁻³)	
						mean	SD	mean	SD
KT4	Kingston Upon Thames - Tolworth Broadway	51.38	-0.28	25.17	laqn	127.00	120.00	51.20	28.10
LB1	Lambeth - Christchurch Road	51.44	-0.12	55.01	laqn	124.00	90.40	57.70	24.90
LB2	Lambeth - Vauxhall Cross	51.48	-0.13	5.44	laqn	131.00	90.10	59.70	22.00
LB4	Lambeth - Brixton Road	51.46	-0.11	12.67	laqn	478.00	298.00	171.00	98.40
LH7	T5 - Oaks Road	51.46	-0.48	22.48	laqn	67.00	84.80	35.50	25.10
LW2	Lewisham - New Cross	51.48	-0.04	13.02	laqn	133.00	106.00	55.40	32.00
LW4	Lewisham - Loampit Vale	51.46	-0.02	7.35	laqn	143.00	118.00	56.10	31.10
MY1	London Marylebone Road	51.52	-0.15	35.00	aun	309.00	219.00	97.30	48.60
MY3	Marylebone Road	51.52	-0.15	27.46	laqn	276.00	169.00	94.20	36.70
NB1	Westminster - Strand (Northbank BID)	51.51	-0.12	14.97	laqn	320.00	207.00	102.00	45.70
NM2	Newham - Cam Road	51.54	-0.00	6.23	laqn	107.00	101.00	51.30	24.20
RB2	Redbridge - Ilford Broadway	51.56	0.07	13.80	laqn	349.00	262.00	122.00	62.50
RB3	Redbridge - Fullwell Cross	51.59	0.09	29.20	laqn	147.00	122.00	60.00	33.60
RB4	Redbridge - Gardner Close	51.58	0.03	23.67	laqn	101.00	103.00	46.50	26.10
RB5	Redbridge - South Woodford	51.60	0.02	44.93	laqn	121.00	99.00	54.50	26.20
RHG	Richmond Upon Thames - Chertsey Road	51.45	-0.34	9.53	laqn	95.00	111.00	39.50	26.00
RI1	Richmond Upon Thames - Castelnau	51.48	-0.24	5.59	laqn	80.90	82.70	40.50	23.80
SK2	Southwark Roadside	51.48	-0.06	10.00	aun	159.00	106.00	62.70	23.20
SK5	Southwark A2 Old Kent Road	51.48	-0.06	10.00	aun	121.00	128.00	49.40	34.00
ST1	Sutton - Robin Hood School	51.37	-0.20	41.85	laqn	108.00	100.00	40.90	24.80
ST4	Sutton - Wallington	51.36	-0.15	61.31	laqn	171.00	141.00	72.90	45.70
ST6	Sutton - Worcester Park	51.38	-0.24	30.84	laqn	136.00	124.00	54.60	33.00
SUT1	Sutton Roadside	51.37	-0.18	40.00	aun	108.00	100.00	41.10	24.90
TH2	Tower Hamlets Roadside	51.52	-0.04	20.00	aun	149.00	117.00	60.80	28.30
TH4	Tower Hamlets - Blackwall	51.52	-0.01	2.88	laqn	155.00	120.00	61.80	27.40
TK2	Thurrock - Purfleet	51.48	0.26	5.84	laqn	192.00	169.00	71.50	35.90
TK8	Thurrock - London Road (Purfleet)	51.48	0.26	6.18	laqn	176.00	166.00	59.90	30.50
VS1	Westminster - Victoria Street	51.50	-0.13	3.80	laqn	189.00	124.00	33.80	17.40
WA4	Wandsworth - High Street	51.46	-0.19	6.50	laqn	98.90	98.80	47.20	27.80
WA7	Wandsworth - Putney High Street	51.46	-0.22	10.23	laqn	330.00	254.00	133.00	89.30
WA8	Wandsworth - Putney High Street Facade	51.46	-0.22	9.73	laqn	255.00	192.00	104.00	64.30
WAA	Wandsworth - Battersea	51.48	-0.14	4.25	laqn	88.30	76.90	41.70	21.90
WAB	Wandsworth - Tooting High Street	51.43	-0.17	16.84	laqn	136.00	101.00	56.80	27.90
WAC	Wandsworth - Lavender Hill (Clapham Jct)	51.46	-0.17	12.53	laqn	108.00	88.50	43.40	24.00
WF1	Watford (Roadside)	51.66	-0.40	76.15	laqn	80.40	79.60	40.20	22.20
WF2	Watford - Watford Town Hall	51.66	-0.40	75.98	laqn	69.10	62.90	37.10	21.40
WL3	Waltham Forest - Chingford	51.63	-0.02	18.40	laqn	63.50	76.80	32.20	20.20
WL4	Waltham Crooked Billet	51.60	-0.02		aqe	193.00	143.00		
WL5	Waltham Forest Leyton	51.56	-0.01		aqe	96.50	118.00		
WM4	Westminster - Charing Cross Library	51.51	-0.13	18.28	laqn	208.00	125.00	84.40	33.00
WM6	Westminster - Oxford Street	51.51	-0.15	24.78	laqn	353.00	266.00	112.00	69.10
ZR2	Dartford Roadside 2 - Town Centre	51.44	0.22	5.64	laqn	103.00	96.10	46.50	27.10
ZY2	Canterbury Roadside - St Dunstons	51.37	-0.29	27.40	laqn	82.90	76.40	36.70	24.70

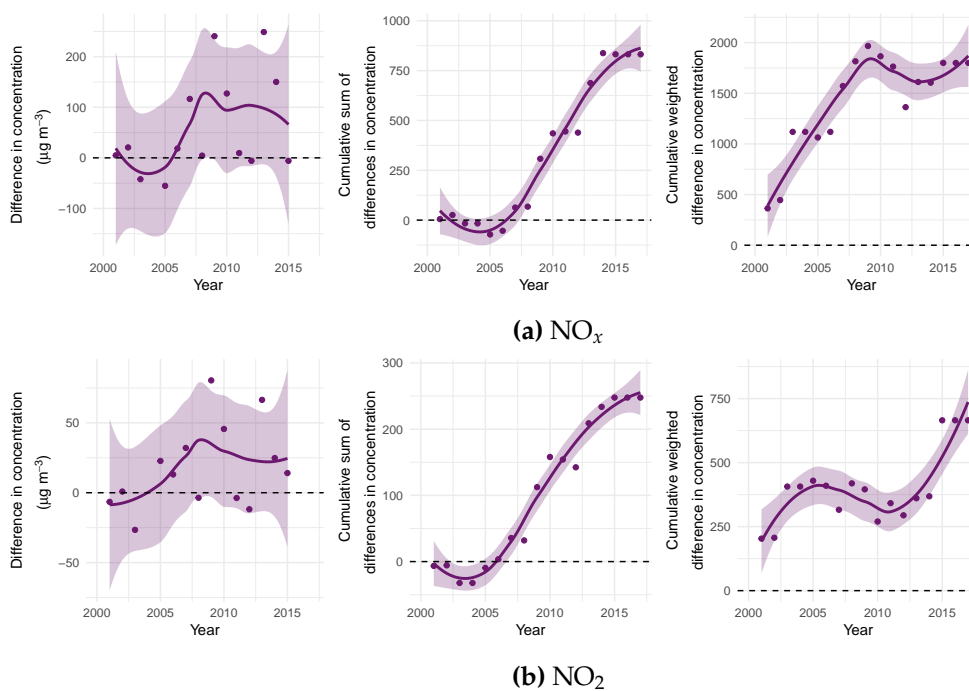
The robustness of the rolling change trend to different values of the moving window width, n was evaluated using data simulated using the methods described in Section 3.1.

Figure B.1 compares the median, 10th worst and worst rolling change trends from a sample of 100 trends to the ‘true trend’ for different values of the window width, n . The percentage errors in the Theil-Sen slope were 26%, 15% and 8% for $n = 2, 3$ and 5 respectively. The accuracy of the rolling change method increases 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.



Figure B.1: Comparison of the rolling change trend of simulated data calculated using window widths $n = 2, 3$, and 5. The trend was calculated from 100 simulated time series, which were randomly sampled 100 times from the ‘combined’ scenario. 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.

The increase in bias in site location towards more polluted sites over time in the London roadside monitoring network 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, as shown in Figure B.2. The difference between the average concentration at sites that are opening and those that are closing is positive (i.e. greater in sites that are opening) over almost all years for NO_x and NO_2 . The cumulative sum of differences in concentration between opening and closing sites demonstrates the compounding effect of the bias over time. Weighting the cumulative sum of differences by the numbers of sites opening and closing gives a further indication of the effect of the bias on the average trend.




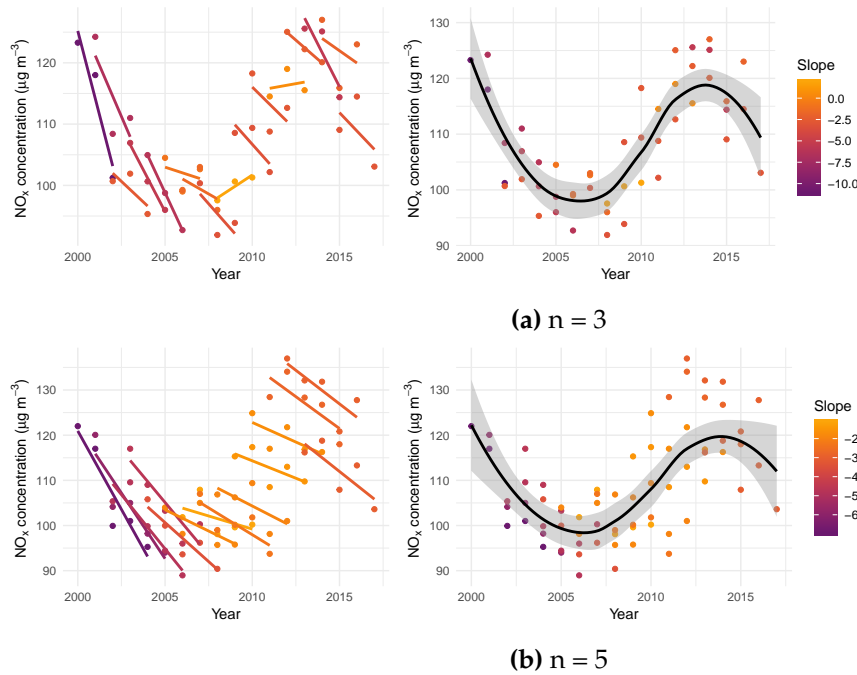


Figure B.2: Difference (left), cumulative sum of differences (centre) and cumulative sum of differences weighted by the number of sites opening and closing (right) in between the average concentration of opening sites and closing sites in each year for (a) NO_x , (b) NO_2 , and (c) NO_2/NO_x at London roadside sites 2000-2017. The lines represent a loess smooth fit to the data, and the shaded bands represent the 95% confidence interval around the smooth.

The effect of the bias in site location on the trends in the average roadside NO_x and NO_2 concentrations, and the average NO_2/NO_x ratio (using data from all monitoring sites) can be observed through comparison of the rolling trends over rolling windows of different widths (n), as shown in Figures B.3, B.4 and B.5.

The effect on the trend in average NO_x concentration (Figure B.3) is particularly clear: the rolling trends reveal that, for small windows which include data from most of the short term sites, an increase in the average trend (right hand plot) is observed between 2008-2013 despite the majority of the rolling trends (left hand plot) having negative gradients. This seems to be an consequence of the changes in magnitude between adjacent rolling trends, presumably due to the inclusion of data from sites opening in the latter years. A similar pattern can be seen in the rolling trends in NO_2 concentration and NO_2/NO_x ratio (Figures B.4 and B.5).



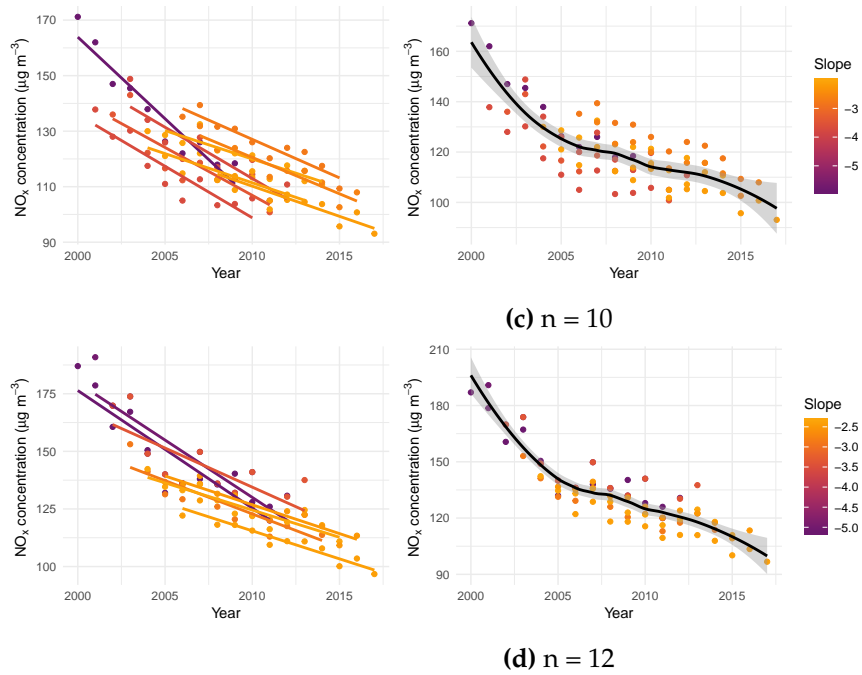
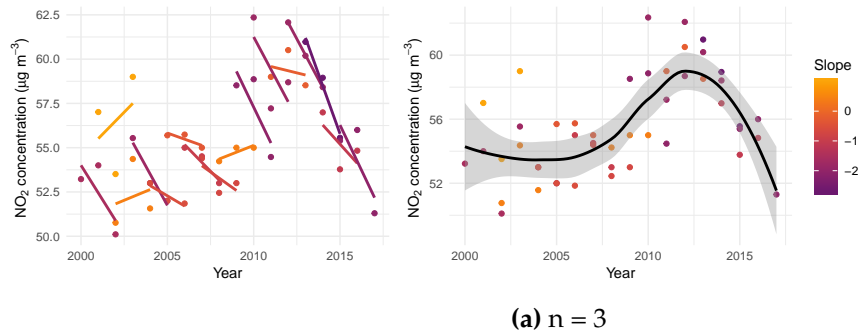


Figure B.3: Rolling trends (left) and average trends (right) in NO_x concentration for $n =$ (a) 3, (b) 5, (c) 10, and (d) 12 at London roadside sites 2000-2017. The average trend (right) was calculated using data from the same sites as the rolling trends, which are filtered by site duration based on the value of n . The black lines in the right hand plots represent a loess smooth fit to the data, and the shaded bands represent the 95% confidence interval around the smooth.



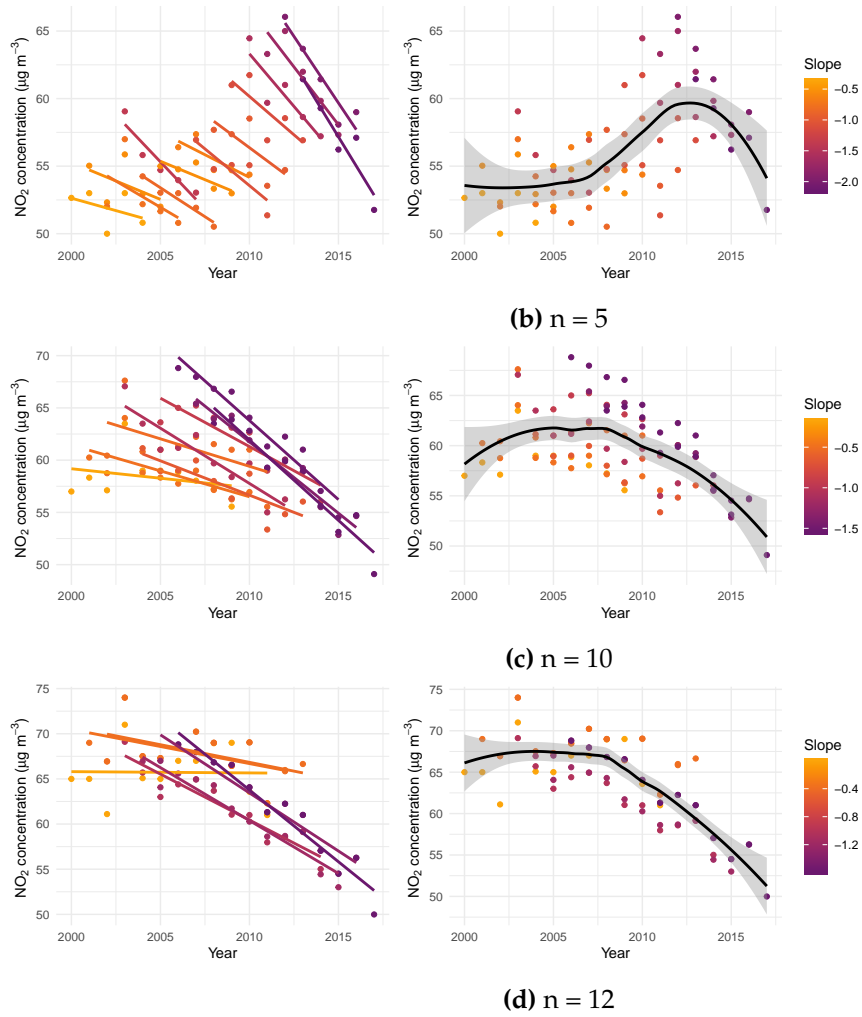


Figure B.4: Rolling trends (left) and average trends (right) in NO_2 concentration for $n =$ (a) 3, (b) 5, (c) 10, and (d) 12 at London roadside sites 2000-2017. The average trend (right) was calculated using data from the same sites as the rolling trends, which are filtered by site duration based on the value of n . The black lines in the right hand plots represent a loess smooth fit to the data, and the shaded bands represent the 95% confidence interval around the smooth.

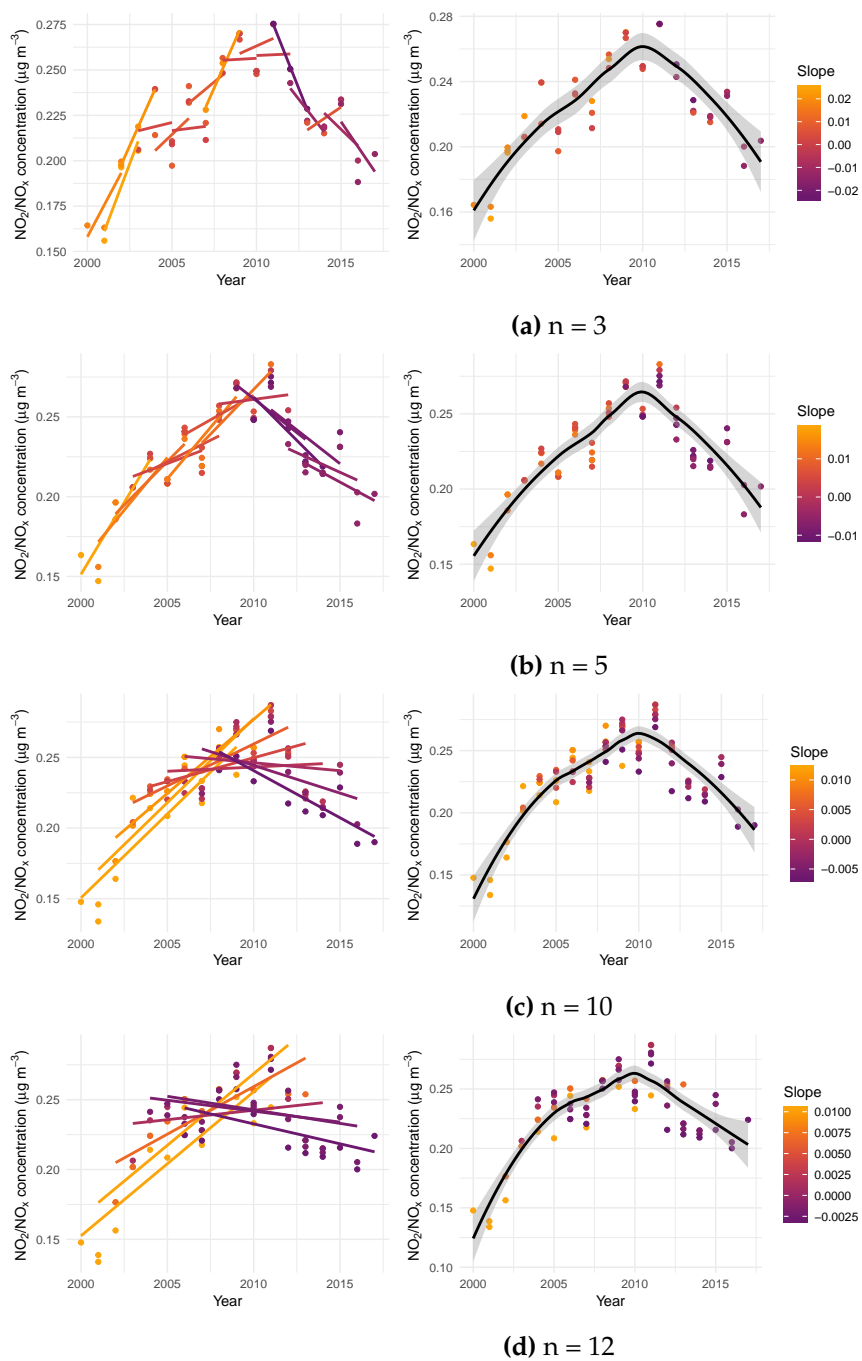


Figure B.5: Rolling trends (left) and average trends (right) in NO_2/NO_x concentration for $n =$ (a) 3, (b) 5, (c) 10, and (d) 12 at London roadside sites 2000-2017. The average trend (right) was calculated using data from the same sites as the rolling trends, which are filtered by site duration based on the value of n . The black lines in the right hand plots represent a loess smooth fit to the data, and the shaded bands represent the 95% confidence interval around the smooth.

Section 3.2 justified the use of the rolling change method for analysis of long term trends in NO_x concentration, NO_2 concentration and the NO_2/NO_x ratio across the London roadside monitoring sites 2000-2017 on the basis that a bias was present in the monitoring networks that influenced the average concentrations. The presence of this bias was demonstrated by comparison of the average long term trend over all sites and over long term sites only, examination of the differences in concentration between opening and closing sites over time, and through the use of rolling trends.

Figures B.6 - B.8 repeat this analysis using data from the London urban background monitoring sites in order to show that the bias does not have as large an effect on these sites, and therefore the rolling change method is unnecessary for analysis of the long term trend.




Figure B.6: Trends in NO_x at London urban background sites 2000-2017 using data from all available monitoring sites (left) and data from long term monitoring sites only (right). 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 average.


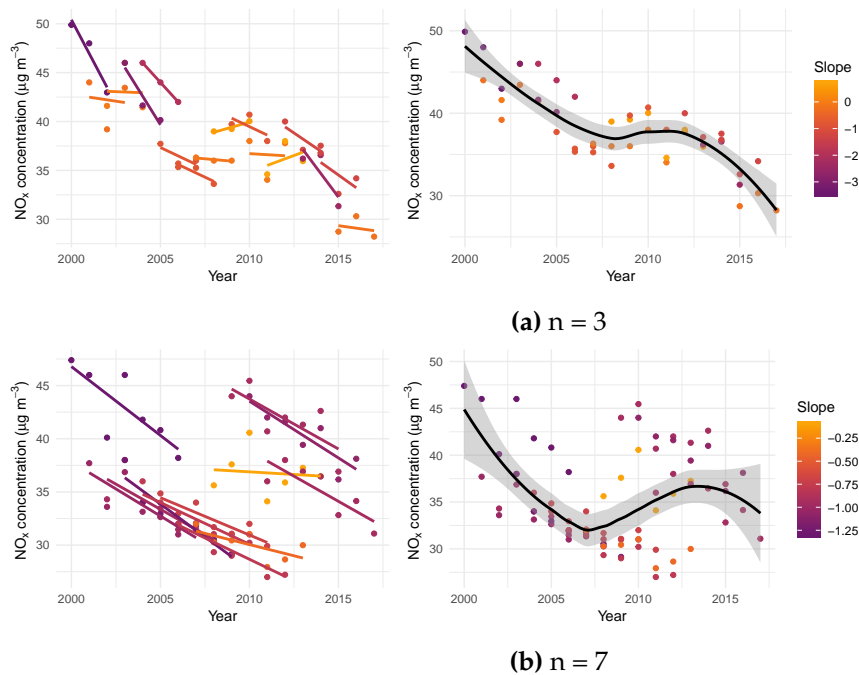


Figure B.7: Difference (left), cumulative sum of differences (centre) and cumulative sum of differences weighted by the number of sites opening and closing (right) in between the average concentration of opening sites and closing sites in each year for NO_x at London urban background sites 2000-2017. The lines represent a loess smooth fit to the data, and the shaded bands represent the 95% confidence interval around the smooth.





(c) $n = 12$

Figure B.8: Rolling trends (left) and average trends (right) in NO_x concentration for $n =$ (a) 3, (b) 7, and (c) 12 at London urban background sites 2000-2017. The average trend (right) was calculated using data from the same sites as the rolling trends, which are filtered by site duration based on the value of n . The black lines in the right hand plots represent a loess smooth fit to the data, and the shaded bands represent the 95% confidence interval around the smooth.

C. Appendix III: aqtrends: An open source R package for air quality trend analysis

aqtrends is an open source R package containing tools for trend analysis of time series data, where the time series are of different lengths. It can be downloaded from the following GitHub page:

<https://github.com/pollylang/aqtrends>

The page also provides a tutorial explaining how to use the functions within the package, illustrated with an example of the usage using London air quality monitoring data. This is shown below.

APPENDIX C. APPENDIX III: AQTRENDS: AN OPEN SOURCE R PACKAGE FOR AIR QUALITY TREND ANALYSIS

16/07/2020

GitHub - pollylang/aqtrends: Tools for trend analysis of air quality network data

 [pollylang](#) / [aqtrends](#)

Tools for trend analysis of air quality network data

☆ 0 stars 🍴 0 forks

☆ Star

👁 Watch

Code

Issues

Pull requests

Actions

Projects

Security

Insights

Join GitHub today

Dismiss

GitHub is home to over 50 million developers working together to host and review code, manage projects, and build software together.

Sign up

🔗 master ▾

 **Polly** committed on 11 Mar 2019



[View code](#)

README.md

Introduction

aqtrends is an R package for conducting trend analysis of time series with different lengths.

It is often useful, in air quality as well as in other fields, to aggregate multiple time series into a single representative trend. In air quality monitoring networks, simply averaging the individual time series often misrepresents the true trend due to biases in the monitoring network. Movement of roadside monitoring sites to more polluted locations is common, as a consequence of legislative requirements to monitor air quality at the most polluted locations. Consequently, opening sites (i.e. starting time series) in later years can leverage the average trend upwards, causing an increase in the average trend that is not necessarily a reflection of trends of individual time series.

<https://github.com/pollylang/aqtrends>

1/7

16/07/2020

GitHub - pollylang/aqtrends: Tools for trend analysis of air quality network data

The `aqtrends` functions are designed to identify and mitigate the biasing effects of monitoring site flux on trends in air pollutant concentrations in air quality monitoring networks.

More details on the methods applied in this package, as well as a detailed case study using London air quality monitoring data, are given in a paper which is currently in preparation.

Installation

To install `aqtrends` from GitHub, the `devtools` package must first be installed. Then copy the following code into R:

```
# Load devtools package
library(devtools)

# Install aqtrends from GitHub
install_github("pollylang/aqtrends")
```

How to use aqtrends

1. Identify the effect of potential biases on the long term trend.
 - Use `average_trends` to compare the average trends of all time series with the average trend for the long term time series only (i.e. time series with equal duration). Also returns average trends of individual sites which can be examined and compared to the average trend to identify biases.
 - Use `site_flux_bias` to visualise the bias due to differences in pollutant concentration between time series that are starting and time series that are ending as a function of time.
 - Use `rolling_trends` to compare rolling trends in pollutant concentration with the average trend.
2. Remove the effect of variable time series length and bias on the long term trend
 - Use `rolling_change_trend` to compute the rolling change trend, which represents the true trend in pollutant concentration.

Example

<https://github.com/pollylang/aqtrends>

2/7

APPENDIX C. APPENDIX III: AQTRENDS: AN OPEN SOURCE R PACKAGE FOR AIR QUALITY TREND ANALYSIS


16/07/2020

GitHub - pollylang/aqtrends: Tools for trend analysis of air quality network data

The `aqtrends` functions are demonstrated below by a trend analysis of monthly NO_x concentration data at 115 roadside monitoring sites in London between 2000-2017. During this period, 9 monitoring sites were constantly open over the duration (long term sites). The data were sourced from [the Automatic Urban and Rural Network](#), [the London Air Quality Network](#) and [Air Quality England](#).

The differences between the average trend (left) and the long term trend (right) in NO_x , plotted using the `average_trends` function, suggest a biasing effect may be influencing the average trend.

```
nox.av.trends <- average_trends(london_nox_data, pollutant = "nox", stat = "me  
start.date = "2000-01-01", end.date = "2017-12  
  
cowplot::plot_grid(nox.av.trends$average.trend, nox.av.trends$longterm.trend,
```




The presence of a biasing effect can be confirmed by plotting the differences in concentration between opening sites and closing sites as a function of year using the `site_flux_bias` function (see below). In this case, it is clear that opening sites have a consistently higher average NO_x concentration than closing sites. Taking into account the relative frequency of opening and closing sites (as shown in the plot of the weighted cumulative sum of differences in concentration as a function of time on the far right), a bias towards more polluted locations is evident.

```
site_flux_bias(london_nox_data, pollutant = "nox", stat = "median")$all
```

APPENDIX C. APPENDIX III: AQTRENDS: AN OPEN SOURCE R PACKAGE FOR AIR QUALITY TREND ANALYSIS


16/07/2020

GitHub - pollylang/aqtrends: Tools for trend analysis of air quality network data



Further evidence of the effect of the bias on the average trend can be visualised using the `rolling_trends` function, as demonstrated below. The plots on the left shows rolling trends over a short moving window, each offset from its neighbours by a single year. The plot on the right shows the average trend over all data included in the rolling trend plots. The larger the width of the moving window, the more constraining the data capture filters on the data. Comparison of the rolling trends and average trends over different moving window widths (i.e. data capture filters - in this case moving window widths = 2, 5, 7, 10, 12, 15 years) demonstrates that some of the features of the average trend, most notably in this case the increase in concentration between 2008-2013, are artefacts of bias in the monitoring network rather than features of the true trend.

```
rolling_trends(london_nox_data, pollutant = "nox", window.width = c(2, 5, 7, 1
## $`moving_window_width=2`
```



```
##
## $`moving_window_width=5`
```

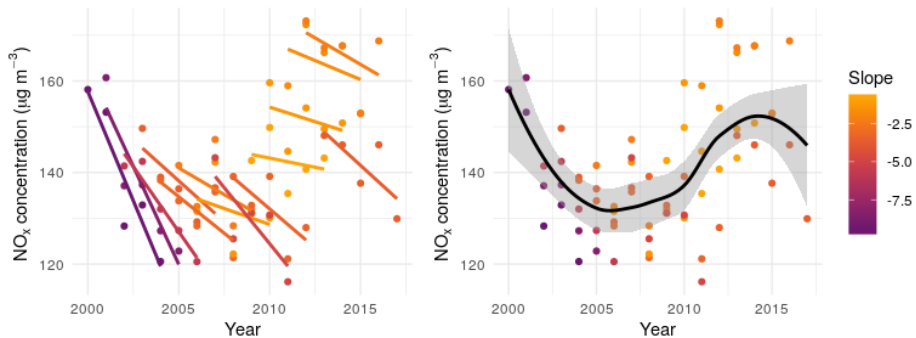
<https://github.com/pollylang/aqtrends>

4/7

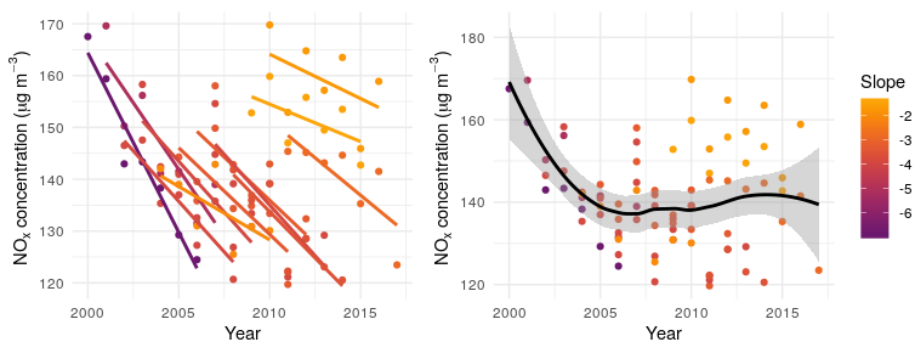
APPENDIX C. APPENDIX III: AQTRENDS: AN OPEN SOURCE R PACKAGE FOR AIR QUALITY TREND ANALYSIS

16/07/2020


GitHub - pollylang/aqtrends: Tools for trend analysis of air quality network data



```
##  
## `$moving_window_width=7`
```



```
##  
## `$moving_window_width=10`
```



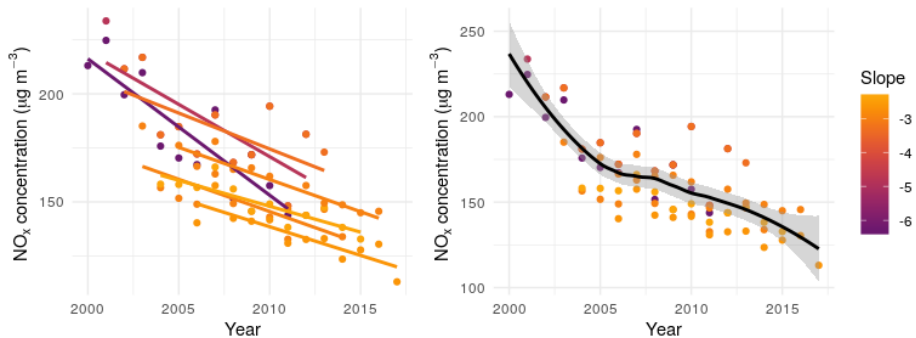
```
##  
## `$moving_window_width=12`
```

<https://github.com/pollylang/aqtrends>

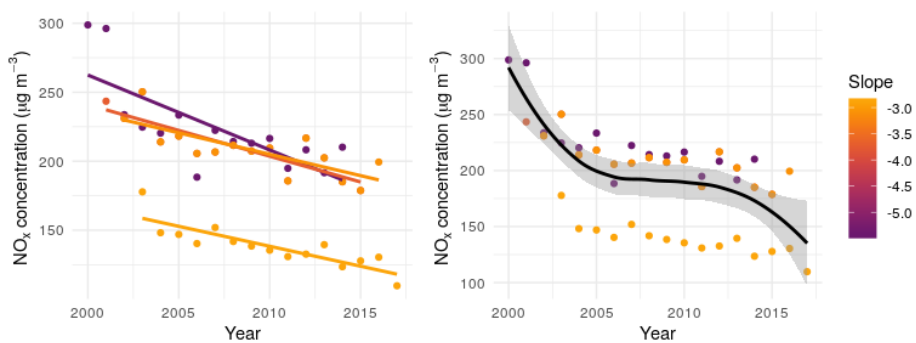
5/7

16/07/2020

GitHub - pollylang/aqtrends: Tools for trend analysis of air quality network data



```
##
## `$moving_window_width=15`
```




Having identified the effect of bias on the average trend, the `rolling_change_trend` function can be used to extract the true trend from the data, as shown below.

```
rolling_change_trend(london_nox_data, "nox", window.width = 3, avg.ts = "year")
```



16/07/2020

GitHub - pollylang/aqtrends: Tools for trend analysis of air quality network data



Releases

No releases published

Languages

● R 100.0%

<https://github.com/pollylang/aqtrends>

7/7

The code for the `rolling_change_method` function is shown below.

APPENDIX C. APPENDIX III: AQTRENDS: AN OPEN SOURCE R PACKAGE FOR AIR QUALITY TREND ANALYSIS

16/07/2020 aqtrends/rolling_change_trend.R at master · pollylang/aqtrends · GitHub

[pollylang / aqtrends](#)

[Code](#) [Issues](#) [Pull requests](#) [Actions](#) [Projects](#) [Security](#) [Insights](#)

Join GitHub today Dismiss

GitHub is home to over 50 million developers working together to host and review code, manage projects, and build software together.

[Sign up](#)

master [aqtrends / R / rolling_change_trend.R](#) [Go to file](#) [...](#)

Polly correct rolling change trend 'date' variable to format YYYY-01-01 Latest commit b442722 on 11 Mar 2019 [History](#)

[1 contributor](#)

182 lines (152 sloc) 7.7 KB [Raw](#) [Blame](#) [Edit](#) [Copy](#)

```
1  #' @title Rolling change trends (extract true trend)
2  #'
3  #' @description Removal of distorting effect of site movement to reveal underlying trend by using changes in
4  #' concentration as a function of year as a proxy for the average trend. This method retains information about the
5  #' shape of the trend, while ignoring differences in magnitude, thus removing the leveraging effect of opening and
6  #' closing sites with extreme magnitudes.
7  #'
8  #' @param obs A data frame of ambient pollutant concentration data. Must contain the columns: site_code, date,
9  #' value. If 'pollutant' is a pollutant ratio, the data frames of the corresponding pollutants
10 #' must be supplied as a list of data frames in the order they are given in the ratio. E.g. for
11 #' \code{pollutant = "no2/nox"}, \code{obs = list(obs.no2, obs.nox)}.
12 #'
13 #' @param pollutant The pollutant of interest (character string). To calculate rolling change trend for a pollutant ratio,
14 #' separate the two pollutants with a forward slash e.g. \code{pollutant = "no2/nox"}.
15 #'
16 #' @param window.width The width of the moving window, n, over which the change in concentration is calculated (in years).
17 #'
18 #' @param avg.ts The resolution to which to average each time series, upon which the rolling regression is carried out.
19 #' For example, setting \code{avg.ts = "day"} means the rolling regression will be carried out on the daily average concentrations
20 #' from each time series (monitoring site). Options are: "year", "month", "week", and "day".
21 #'
22 #' @param stat The metric (character string) used to average the ambient concentration data by year. Options: "median", "mean".
23 #'
24 #' @param start.date,end.date The starting and ending dates (character string) of the period of interest over which
25 #' to calculate and plot the change trend.
26 #'
27 #' @param parallel Logical indicating whether the rolling changes should be computed in parallel. If \code{TRUE}, the
28 #' parallelisation will be implemented using the \code{foreach} function. The number of cores used will be the total
29 #' number of cores - 1.
30 #'
31 #' @param verbose Logical indicating whether to print the date range of the rolling window over which the calculation is
32 #' being applied.
33 #'
34 #' @return A plot of the rolling change trend.
35 #'
36 #' @import dplyr
37 #'
38 #' @examples
39 #' \dontrun{
40 #'   rolling_change_trend(london_nox_data,
41 #'     pollutant = "nox",
42 #'     window.width = 3,
43 #'     avg.ts = "year",
44 #'     stat = "median",
45 #'     start.date = "2000-01-01", end.date = "2017-12-31")
46 #' }
47 #'
```

https://github.com/pollylang/aqtrends/blob/master/R/rolling_change_trend.R 1/3

APPENDIX C. APPENDIX III: AQTRENDS: AN OPEN SOURCE R PACKAGE FOR AIR QUALITY TREND ANALYSIS

```
16/07/2020 aqtrends/rolling_change_trend.R at master · pollylang/aqtrends · GitHub
48 #' @export
49
50
51 rolling_change_trend <- function(obs,
52                                 pollutant,
53                                 window.width,
54                                 avg.ts = "year",
55                                 stat = "median",
56                                 start.date = "2000-01-01",
57                                 end.date = "2017-12-31",
58                                 parallel = FALSE,
59                                 verbose = FALSE){
60
61
62 ## Check arguments
63 check_arguments(obs = obs,
64                 pollutant = pollutant,
65                 window.width = window.width,
66                 stat = stat,
67                 start.date = start.date,
68                 end.date = end.date,
69                 parallel = parallel,
70                 avg.ts = avg.ts)
71
72
73 # Function to calculate data frame with row for difference between initial conc and trend for each moving window
74 trend_difference <- function(rolling.x){
75
76 # Initialise delta_y1 (concentration change value in the first year = raw annual average concentration in first year)
77 first.row <- rolling.x[1, ] %>% dplyr::mutate(trend = av_value) # copy the first row to the data frame (initialise delta y1)
78
79 rolling.x <- rolling.x %>%
80   tibble::add_row(date = first.row$date, av_value = first.row$av_value, n = first.row$n,
81                 moving_window = paste0(first.row$moving_window, ".1"), window_width = first.row$window_width,
82                 trend = first.row$trend, .before = 1) %>% # add first row (initialise - delta y1)
83   dplyr::mutate(moving_window = ifelse(moving_window == first.row$moving_window,
84                                     paste0(first.row$moving_window, ".2"),
85                                     moving_window)) %>%
86   dplyr::group_by(moving_window) %>%
87   dplyr::summarise(trend = unique(trend),
88                  n = unique(n),
89                  date = (min(date) + floor((max(date)-min(date))/2)),
90                  window_width = unique(window_width)) %>%
91   dplyr::ungroup() %>%
92   dplyr::arrange(date) %>%
93   dplyr::mutate(trend_difference = ifelse(moving_window == paste0(first.row$moving_window, ".1"),
94                                         trend, NA))
95
96 # Calculate concentration change: delta y_i = delta y_{i-1} + beta_i
97 for(i in 2:nrow(rolling.x)){
98
99   rolling.x[i, "trend_difference"] <- rolling.x[i-1, "trend_difference"] + rolling.x[i, "trend"]
100
101 }
102
103 return(rolling.x)
104 }
105
106
107 # Plot rolling change trend (concentration change as a function of the year)
108 period_plot <- function(df){
109
110 plot.colour <- viridis::inferno(1, begin=0.3, end=0.8)
111
112 p <- df %>%
113   ggplot2::ggplot(ggplot2::aes(x = date, y = trend_difference)) +
114   ggplot2::geom_point(color = plot.colour) +
115   ggplot2::geom_smooth(method = "loess",
116                       color = plot.colour,
117                       fill = plot.colour,
118                       alpha=0.25) +
119   ggplot2::xlab("Year") +
120   ggplot2::ylab(openair::quickText(paste0(pollutant, " concentration (ug m-3)")) +
121   ggplot2::geom_text(ggplot2::aes(label = n), vjust = 1.3, size = 3) +
122   ggplot2::theme_minimal()
123
124
125
126
127
128
129
130
131
132
133
134
135
136
137
138
139
140
141
142
143
144
145
146
147
148
149
150
151
152
153
154
155
156
157
158
159
160
161
162
163
164
165
166
167
168
169
170
171
172
173
174
175
176
177
178
179
180
181
182
183
184
185
186
187
188
189
190
191
192
193
194
195
196
197
198
199
200
201
202
203
204
205
206
207
208
209
210
211
212
213
214
215
216
217
218
219
220
221
222
223
224
225
226
227
228
229
230
231
232
233
234
235
236
237
238
239
240
241
242
243
244
245
246
247
248
249
250
251
252
253
254
255
256
257
258
259
260
261
262
263
264
265
266
267
268
269
270
271
272
273
274
275
276
277
278
279
280
281
282
283
284
285
286
287
288
289
290
291
292
293
294
295
296
297
298
299
300
301
302
303
304
305
306
307
308
309
310
311
312
313
314
315
316
317
318
319
320
321
322
323
324
325
326
327
328
329
330
331
332
333
334
335
336
337
338
339
340
341
342
343
344
345
346
347
348
349
350
351
352
353
354
355
356
357
358
359
360
361
362
363
364
365
366
367
368
369
370
371
372
373
374
375
376
377
378
379
380
381
382
383
384
385
386
387
388
389
390
391
392
393
394
395
396
397
398
399
400
401
402
403
404
405
406
407
408
409
410
411
412
413
414
415
416
417
418
419
420
421
422
423
424
425
426
427
428
429
430
431
432
433
434
435
436
437
438
439
440
441
442
443
444
445
446
447
448
449
450
451
452
453
454
455
456
457
458
459
460
461
462
463
464
465
466
467
468
469
470
471
472
473
474
475
476
477
478
479
480
481
482
483
484
485
486
487
488
489
490
491
492
493
494
495
496
497
498
499
500
501
502
503
504
505
506
507
508
509
510
511
512
513
514
515
516
517
518
519
520
521
522
523
524
525
526
527
528
529
530
531
532
533
534
535
536
537
538
539
540
541
542
543
544
545
546
547
548
549
550
551
552
553
554
555
556
557
558
559
560
561
562
563
564
565
566
567
568
569
570
571
572
573
574
575
576
577
578
579
580
581
582
583
584
585
586
587
588
589
590
591
592
593
594
595
596
597
598
599
600
601
602
603
604
605
606
607
608
609
610
611
612
613
614
615
616
617
618
619
620
621
622
623
624
625
626
627
628
629
630
631
632
633
634
635
636
637
638
639
640
641
642
643
644
645
646
647
648
649
650
651
652
653
654
655
656
657
658
659
660
661
662
663
664
665
666
667
668
669
670
671
672
673
674
675
676
677
678
679
680
681
682
683
684
685
686
687
688
689
690
691
692
693
694
695
696
697
698
699
700
701
702
703
704
705
706
707
708
709
710
711
712
713
714
715
716
717
718
719
720
721
722
723
724
725
726
727
728
729
730
731
732
733
734
735
736
737
738
739
740
741
742
743
744
745
746
747
748
749
750
751
752
753
754
755
756
757
758
759
760
761
762
763
764
765
766
767
768
769
770
771
772
773
774
775
776
777
778
779
780
781
782
783
784
785
786
787
788
789
790
791
792
793
794
795
796
797
798
799
800
801
802
803
804
805
806
807
808
809
810
811
812
813
814
815
816
817
818
819
820
821
822
823
824
825
826
827
828
829
830
831
832
833
834
835
836
837
838
839
840
841
842
843
844
845
846
847
848
849
850
851
852
853
854
855
856
857
858
859
860
861
862
863
864
865
866
867
868
869
870
871
872
873
874
875
876
877
878
879
880
881
882
883
884
885
886
887
888
889
890
891
892
893
894
895
896
897
898
899
900
901
902
903
904
905
906
907
908
909
910
911
912
913
914
915
916
917
918
919
920
921
922
923
924
925
926
927
928
929
930
931
932
933
934
935
936
937
938
939
940
941
942
943
944
945
946
947
948
949
950
951
952
953
954
955
956
957
958
959
960
961
962
963
964
965
966
967
968
969
970
971
972
973
974
975
976
977
978
979
980
981
982
983
984
985
986
987
988
989
990
991
992
993
994
995
996
997
998
999
1000
```

https://github.com/pollylang/aqtrends/blob/master/R/rolling_change_trend.R

2/3

APPENDIX C. APPENDIX III: AQTRENDS: AN OPEN SOURCE R PACKAGE FOR AIR QUALITY TREND ANALYSIS

```
16/07/2020 aqtrends/rolling_change_trend.R at master · pollylang/aqtrends · GitHub
123
124     return(p)
125 }
126 }
127
128 # Define the start and end dates
129 window1 <- define_moving_windows(window.width, obs, start.date, end.date, avg.ts)
130
131 # For each moving window, compute rolling regression
132 if(parallel == TRUE){
133   no_cores <- parallel::detectCores() - 1
134   cl <- parallel::makeCluster(no_cores)
135   doParallel::registerDoParallel(cl)
136
137   rolling.df <- foreach::foreach(d1 = window1, .combine = "rbind", .packages = c("dplyr", "ggplot2", "lubridate"),
138     .export = c("average_data", "sites_open_throughout_window", "calculate_pollutant_ratio", "rolling_worker"),
139     try(rolling_worker(d1, obs = obs, pollutant = pollutant, window.width = window.width, stat = stat, avg.ts = avg.ts,
140       verbose = verbose))
141
142   doParallel::stopImplicitCluster()
143 } else{
144   rolling.df <- tryCatch({purrr::map_dfr(window1, rolling_worker,
145     obs = obs,
146     pollutant = pollutant,
147     window.width = window.width,
148     stat = stat, avg.ts = avg.ts,
149     verbose = verbose)
150 }, error = function(e){
151   print("Error message: ")
152   print(e)
153   data.frame("date" = as.POSIXct(character()),
154     "av_value" = numeric(),
155     "n" = numeric())
156 })
157 }
158
159
160 if(nrow(rolling.df) < 1){
161   print("Error in rolling trend calculation. Possibly insufficient data.")
162   return(NULL)
163 } else{
164
165   # Calculate concentration change for each year (moving window)
166   rolling.out <- trend_difference(rolling.df)
167
168   # Correct the 'date' variable to always be in the form YYYY-01-01 (where YYYY is the middle year of the moving window)
169   rolling.out <- rolling.out %>%
170     mutate(year = case_when(lubridate::month(date) == 1 ~ lubridate::year(date),
171       lubridate::month(date) == 12 ~ lubridate::year(date)+1,
172       TRUE ~ lubridate::year(date))) %>%
173     mutate(date = lubridate::ymd(paste0(year, "-01-01")))
174
175   # Plot rolling change trend
176   plots <- period_plot(rolling.out)
177
178   return(plots)
179 }
180 }
181 }
182 }
```

D. Appendix IV: Publication

This appendix contains the paper presented alongside a talk at the 2019 International Transport and Air Pollution Conference.

Lang, P. E., Carslaw, D. C., & Moller, S. J. (2019). Analysis of long term roadside air pollution trends in European cities using ambient data from sparse monitoring networks. In 23rd International Transport and Air Pollution Conference.

Analysis of long term roadside air pollution trends in European cities using ambient data from sparse monitoring networks

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

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

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

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

1. Introduction

2 Air pollution is one of the most important problems facing Europe, re-
3 sponsible for an estimated 400,000 premature deaths a year (EEA, 2018b). Of
4 particular concern are the concentrations of nitrogen oxides (NO_x), composed
5 of NO and NO₂, and particulate matter (PM).

6 Frequent and widespread exceedances of the European Union limits
7 on NO₂ and PM concentration have motivated policies aimed at reducing
8 the concentration of these air pollutants. These policies have included the
9 reduction of vehicle emissions through European Directives (Euro Standards)
10 that have set increasingly stringent limits on the emissions from road vehicles.
11 To meet these emission limits, vehicle manufacturers have adopted exhaust
12 technologies such as three-way catalysts on petrol vehicles, diesel particulate
13 filters (DPF) and diesel oxidation catalysts (DOC) (EEA, 2016). The impact
14 of these changes on ambient air quality in Europe is difficult to establish
15 but can potentially be evaluated through analysis of the long term trends
16 in ambient roadside concentrations of air pollutants using data from the
17 extensive European monitoring network. Such an evaluation also has the
18 potential to provide information on the effectiveness of different policies to
19 reduce air pollution.

20 The vast amount of ambient monitoring data measured at thousands of
21 monitoring stations across Europe provides the ability to directly analyse

*Corresponding author
Email address: p1746@york.ac.uk (Polly E. Lang)

22 the large-scale changes in roadside air quality in Europe over time. While
23 the data from a single monitoring station provides information about the
24 local variation in air quality at a specific location, aggregation of data from
25 multiple monitoring stations enables the effects of local variability to be
26 averaged out, leaving a better indication of the large-scale trend.

27 However, there are several limitations inherent in the established methods
28 for calculating aggregate trends using data from air quality monitoring
29 networks. One approach is to compare individual time series from different
30 monitoring stations (e.g. Masiol et al. (2017); Mavroidis and Chaloulakou
31 (2011)), however this becomes impractical in very large monitoring networks,
32 such as the European network.

33 Another method involves evaluating the trend in the average concentration
34 across all the available monitoring sites. This trend can be biased due to the
35 leveraging effects of site flux (i.e. sites opening and closing during the period
36 of interest) on the average concentration, as demonstrated by Lang et al.
37 (2019) in a trend analysis using data from the London monitoring network.
38 For example, in a given year, the opening or closing of monitoring stations
39 cause new data to be included in the average concentration, potentially
40 resulting in abrupt changes in the average concentration that are driven,
41 not by changes in source emissions strength, but by the sudden inclusion or
42 exclusion of data. Many monitoring networks, particularly those composed
43 of roadside stations, experience considerable site flux, which can have a
44 substantial effect on the calculated trend in average concentration.

45 To mitigate this issue, data filtering can be applied to ensure that only
46 data from monitoring stations with a complete time series (i.e. measuring
47 constantly over the period of the trend analysis) are included in the analysis
48 (e.g. Font and Fuller (2016)). This approach inevitably results in the
49 exclusion of much of the available data and therefore a considerable loss of
50 information. The rolling change method is a technique designed to address
51 the above issues, and enable the calculation of aggregate trends in ambient
52 air quality using data from monitoring networks that evolve over time. The
53 method is robust to the biasing effects of site flux, while preserving more
54 data in the trend analysis than data filtering methods (Lang et al., 2019).
55 This approach enables better investigation of the large scale air pollution
56 trends across an area of interest providing the ability to produce trends that
57 are more representative, due to the inclusion of more sites, and driven by
58 genuine concentration changes.

59 This study uses the rolling change method to calculate large-scale trends
60 in roadside concentrations of NO_x , NO_2 , PM_{10} and $\text{PM}_{2.5}$ in Europe between
61 2000 and 2017. The trends in roadside air quality are also calculated in 44

62 individual European functional urban areas (FUAs), and compared with the
63 European-wide trends. The reasons for the observed changes in air pollutant
64 concentrations are discussed in the context of changes in vehicle emissions,
65 and possible explanations for the deviations from the overall trend pattern
66 observed in some FUAs are proposed.

67 **2. Methods and Data**

68 The long term trends for Europe and for each European FUA were
69 calculated using the rolling change method (Lang et al., 2019). The method
70 involves splitting the period of interest into short moving windows, each
71 offset from the adjacent window by a time step, in this case one year. The
72 data within each window is filtered to exclude data from monitoring sites
73 with an incomplete time series over the period of the moving window (i.e.
74 sites which open or close during the window), to ensure that all time series
75 within the moving window are the same length. The change in concentration
76 over the window is calculated by fitting a linear regression to this filtered
77 data. By shifting the moving window along the period of analysis, and
78 calculating the concentration change at each step, a trend can be estimated.
79 This aggregate trend represents the average change in concentration over
80 all monitoring sites, unaffected by the leveraging effects of changes in the
81 average concentration caused by site flux during the period of analysis. In
82 effect, the method relaxes the data capture requirement to short time periods
83 rather than the whole time series, which increases the amount of data that
84 can be used in trend analysis.

85 The rolling change method was applied using annual median data, re-
86 sulting in median trends for each FUA. The advantage of an annual trend,
87 as opposed to a trend with more granular resolution, is to ‘average out’
88 the effects of confounders which vary on a shorter time scale, for example,
89 seasonal variation.

90 All ambient concentration data were obtained from the smonitor Europe
91 database (Grange, 2016), which stores data collected from the European
92 Environment Agency AirBase and air quality e-Reporting data repositories
93 (EEA, 2018a, 2019). Only data from roadside monitoring sites were con-
94 sidered. The distribution of monitoring sites contributing to the aggregate
95 trends over the whole of Europe is displayed in Figures 1-2. In each case,
96 the ‘number of sites’ corresponds to the total number of unique sites that
97 contributed to the trend over the entire period of analysis (2000-2017).

98 Figures 1 and 2 indicate that the distribution of monitoring sites in the
99 European network is highly heterogeneous, with several countries (Italy,

100 Germany, Spain, France and the UK) dominating the network. Consequently,
 101 these countries will dominate the shape of the overall European trends. This
 102 provides an additional motivation for examining the trends at a city-scale
 103 in complement to the European-scale, as it allows closer analysis of regions
 104 which are underrepresented in the European-wide trend.




Figure 1: The total number of unique roadside monitoring sites contributing to the rolling change trends in NO_x, NO₂, PM₁₀ and PM_{2.5} concentration over the period 2000-2017, by country and pollutant.

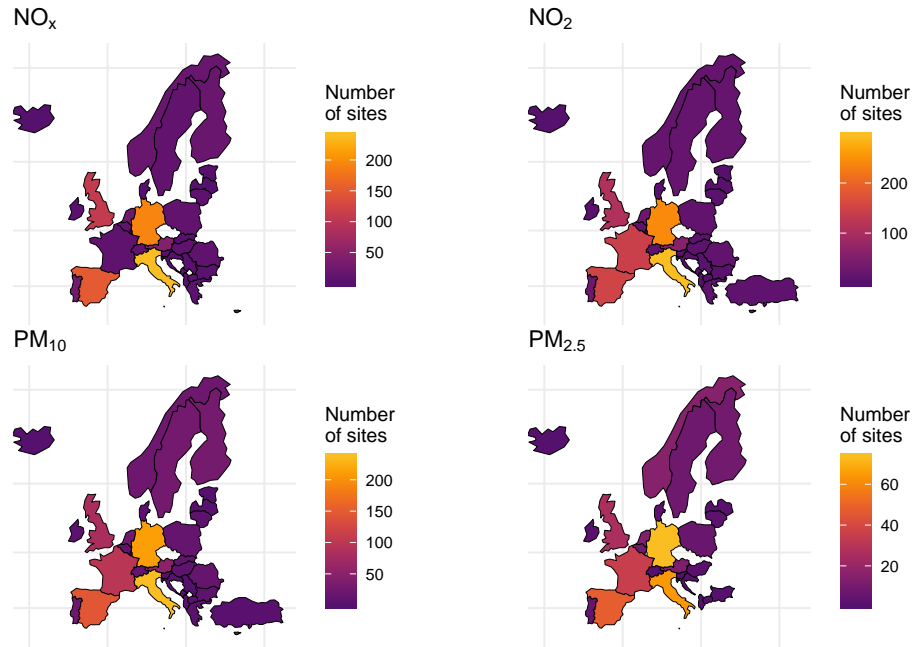


Figure 2: Distribution of roadside monitoring sites contributing to the rolling change trends in NO_x, NO₂, PM₁₀ and PM_{2.5} concentration across Europe. The colours represent the total number of unique roadside monitoring sites contributing to the rolling change trend in each country between 2000 and 2017.

105 The European city in which each monitoring sites was located was
 106 determined using spatial data on the European functional urban areas (FUA).
 107 The FUA spatial data were obtained from the European Commission Joint
 108 Research Centre (JRC) (Lavallo et al., 2015). The criteria for an FUA to
 109 be included in the analysis was that the FUA must (i) contain at least
 110 one monitoring site measuring the pollutant of interest, and (ii) not have
 111 a period equal to or greater than the width of the moving window (three
 112 years) during which there are no measurements of the pollutant of interest
 113 available. The FUAs for which trend analysis was conducted, according to
 114 these criteria, are shown in Table 1, along with the total number of distinct
 115 roadside monitoring sites that contributed to the trend in each FUA (over
 116 the entire time period of the analysis).

117 The analysis was carried out in R, using the aqtrends R package for the
 118 calculation of the rolling change trend (Lang, 2018).

Table 1: Total number of unique roadside monitoring sites measuring NO_x, NO₂, PM₁₀ and PM_{2.5} concentration in each European functional urban area analysed.

FUA	NO ₂	NO _x	PM ₁₀	PM _{2.5}
Amsterdam	4	0	0	0
Athina	3	0	0	0
Barcelona	10	10	13	0
Berlin	8	8	7	3
Bilbao	4	3	0	0
Bologna	3	0	0	0
Bruxelles / Brussel	5	0	0	0
Frankfurt am Main	4	0	3	0
Genova	3	0	0	0
Gijn	2	2	2	0
Hamburg	5	0	3	0
Helsinki	3	3	3	0
Innsbruck	4	0	3	0
Karlsruhe	0	0	3	0
La Spezia	2	0	0	0
Lille	2	0	0	0
Linz	3	0	0	0
Lisboa	3	3	3	0
London	25	25	20	3
Lyon	5	0	3	0
Madrid	11	10	8	3
Mainz	4	4	0	0
Mannheim-Ludwigshafen	6	4	4	2
Marseille	3	0	0	0
Milano	7	0	0	0
Mnchen	3	3	3	0
Nrnberg	3	3	3	0
Paris	6	0	4	0
Pescara	3	0	0	0
Porto	5	5	4	0
Praha	7	7	6	3
Roma	7	7	5	0
Rotterdam	0	0	0	2
Ruhrgebiet	7	4	0	0
Salzburg	4	0	3	0
Stockholm	3	3	4	0
Stuttgart	6	0	5	0
Torino	3	0	0	0

Toulouse	3	0	0	0
Udine	2	0	0	0
Utrecht	3	3	0	0
Valencia	3	3	0	0
Valladolid	0	0	2	0
Wien	12	0	0	0

119 **3. Results and discussion**

120 Figure 3 shows the long term trends in roadside NO_x , NO_2 , PM_{10} and
 121 $\text{PM}_{2.5}$ concentrations, as well as the NO_2/NO_x , $\text{PM}_{10}/\text{NO}_x$ and $\text{PM}_{2.5}/\text{NO}_x$
 122 ratios between 2000 and 2017 across the entirety of Europe. It can be seen
 123 that, on average, in Europe the concentrations of NO_x and $\text{PM}_{2.5}$ have
 124 been decreasing monotonically since 2000. NO_2 and PM_{10} concentrations
 125 increased slightly to 2003, and decreased monotonically subsequently.

126 The observed decrease in NO_x and NO_2 concentrations in Europe since
 127 2000 and 2002 respectively can be attributed to the introduction of vehicle
 128 exhaust technologies aimed at reducing emissions of these species, such as
 129 three-way catalysts used on petrol vehicles and the more recent use of Lean
 130 NO_x Traps (LNT) and Selective Catalytic Reduction (SCR) on diesel vehicles
 131 over this period. Indeed, the rolling trend analysis of roadside NO_x and
 132 NO_2 concentrations provide compelling evidence that emissions of NO_x have
 133 strongly decreased across Europe as a whole. These conclusions are robust in
 134 the sense the rolling trend analysis maximises the use of the large amounts
 135 of measurement data available across Europe while largely eliminating any
 136 bias introduced due to differing numbers of sites available in each year.

137 The increase in NO_2 concentration prior to 2002 was most likely due to
 138 the increase in the number of diesel passenger vehicles over this period. As
 139 shown in Figure 4, the proportion of the new vehicle registrations in Europe
 140 composed of diesel vehicles increased from 36% in 2001 to 48% in 2005, then
 141 increased more slowly to a peak in 2011/12 of 55%, before starting to decrease
 142 (The International Council on Clean Transportation, 2018). This period also
 143 coincided with the introduction of DOCs to new vehicles in compliance with
 144 the Euro 3 and Euro 4 emission standards, leading to the emission of more
 145 NO_2 from vehicle exhaust (Carslaw et al., 2019).

146 A study by Grange et al. (2017) considered the trends in average NO_x and
 147 NO_2 concentrations in Europe, and found that NO_x concentrations decreased
 148 between 1998 and 2015, in corroboration with the findings of this study.
 149 The principal finding of Grange et al. (2017) was that the emissions ratio of
 150 NO_2/NO_x from road vehicles increased from 2000 to about 2009 and then

151 started to decrease. A consideration of the ambient roadside concentration
152 NO_2/NO_x ratios also shows a clear increase from 2000 to about 2009 and
153 then a decreasing trend. This behaviour is entirely consistent with Grange
154 et al. (2017) and is related to increased ratio of NO_2/NO_x in vehicle exhaust
155 due to DOC and DPF. In DOC and DPF, the deliberate production of NO_2
156 through oxidation of NO is used to help oxidise other pollutants such as CO ,
157 hydrocarbons and particulate matter.

158 The more recent decrease in the ambient NO_2/NO_x ratio likely has several
159 origins. Recent vehicle emission remote sensing measurements show that
160 as the vehicle mileage increases for diesel passenger cars, the NO_2/NO_x
161 ratio decreases Carslaw et al. (2019). This deterioration effect is unlikely
162 to be the only factor affecting the recent decreasing NO_2/NO_x ratio. It is
163 also likely that vehicle emission after-treatment systems have been better
164 optimised so as not over-produce NO_2 . These other factors require further
165 emission measurements and analysis to confirm their contribution. However,
166 given the timing of the change shown in Figure 3, there is consistency
167 with the introduction of Euro 5/V Standards around 2009. It should also
168 be noted that as NO_x and NO_2 concentrations continue to decrease, the
169 ambient concentration ratio of NO_2/NO_x will eventually increase due to the
170 increased availability of O_3 to convert NO to NO_2 , which can be readily
171 confirmed by plotting the NO_2/NO_x ratio against NO_x . For example, at
172 urban background sites, the ratio of NO_2/NO_x is typically around 0.6 to 0.7,
173 which is considerably higher than the peak ratio seen in Figure 3. There
174 might be some indication of this increase shown in Figure 3, although more
175 data are required to confirm whether an increasing trend continues. Such
176 an increase will also be dominated by the increased role of O_3 rather than
177 primary NO_2 emissions from vehicles.

178 The trend plots for PM_{10} and $\text{PM}_{2.5}$ both show substantial reductions in
179 concentration since 2000. While the data shown in Figure 3 are for roadside
180 sites, the concentrations of both species will be strongly influenced by non
181 road sources including secondary particulate matter (principally sulphate
182 and nitrate). A more informative analysis is to consider the ratio of PM to
183 NO_x , where NO_x acts as a local combustion tracer. For PM_{10} and $\text{PM}_{2.5}$,
184 the ratio to NO_x has decreased in recent years, which likely reflects the
185 increased use of DPF on heavy and light duty vehicles. The trends in the
186 ratios shown in Figure 3 show that there has been a greater reduction in
187 the ratio of $\text{PM}_{2.5}/\text{NO}_x$ than $\text{PM}_{10}/\text{NO}_x$. This behaviour is consistent with
188 the more effective reduction in $\text{PM}_{2.5}$ exhaust emissions than total PM_{10} ,
189 the latter of which will also include a stronger influence of coarse fraction
190 ($\text{PM}_{2.5}$ to PM_{10}) tyre, brake and road abrasion sources, which have not been

191 mitigated.

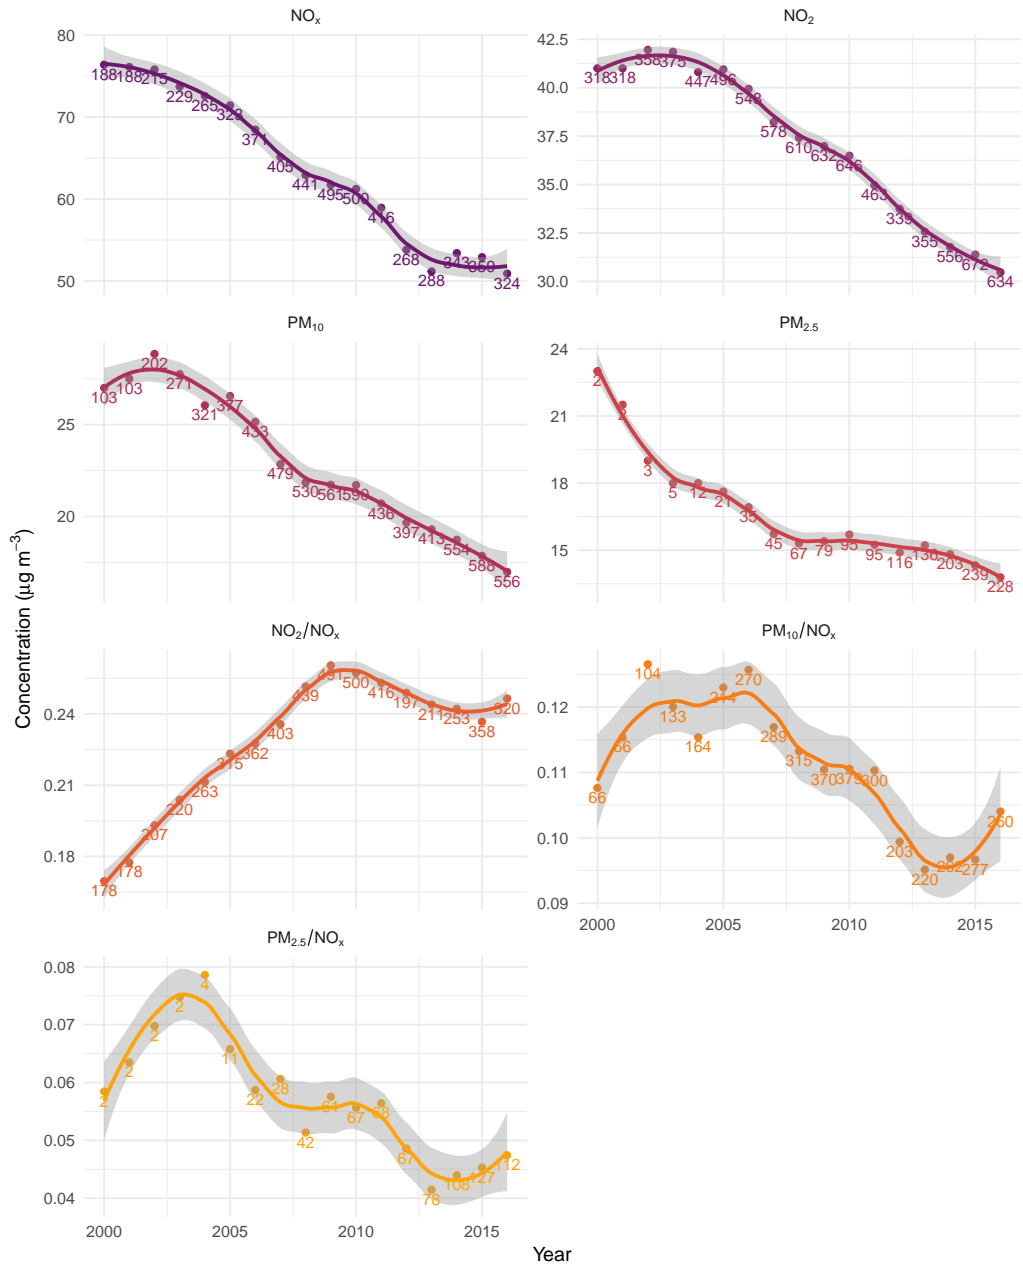


Figure 3: Rolling change trends in NO_x, NO₂, PM₁₀ and PM_{2.5} roadside concentration, and the NO₂/NO_x, PM₁₀/NO_x, and PM_{2.5}/NO_x ratios across Europe between 2000-2017. The smoothed lines are loess (local regression) fits, with the 95% confidence interval represented by the shaded band. The numbers signify the number of monitoring sites contributing to each annual data point. Annual pollutant ratios were calculated as the slope of a regression of the hourly concentration of one pollutant against the other (for each year of data).




Figure 4: Proportion of new registrations composed of diesel vehicles in Europe by year between 2001 and 2017 (The International Council on Clean Transportation, 2018). The smoothed line is a loess (local regression) fit, with the 95% confidence interval represented by the shaded band.

192 The long term trends in the ambient roadside concentrations of NO_x ,
 193 NO_2 , PM_{10} and $\text{PM}_{2.5}$ in individual functional urban areas (FUAs) are shown
 194 in Figure 5. The trends in the majority of FUAs closely resemble the whole-
 195 European trends (Figure 3), however a visual inspection of the individual
 196 trends revealed that some FUAs exhibit trends that differ considerably from
 197 the consensus pattern. These outliers are highlighted in bold in Figure 5,
 198 and shown individually in Figure 6.

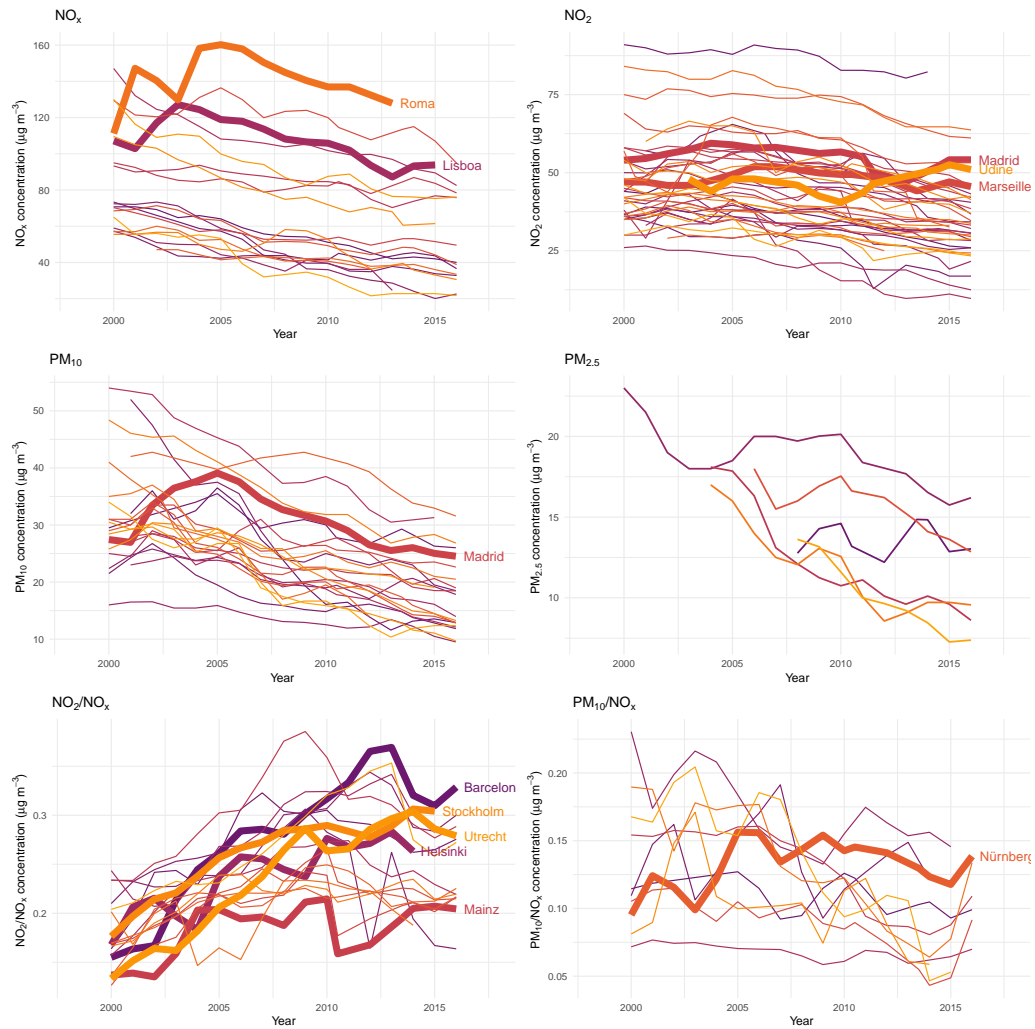


Figure 5: Rolling change trends in NO_x , NO_2 , PM_{10} and $\text{PM}_{2.5}$ roadside concentration in individual European FUAs between 2000-2017. The FUAs with trends differing considerably from the overall European trend are shown in bold and labelled.



Figure 6: Rolling change trends in NO_x , NO_2 , PM_{10} and $\text{PM}_{2.5}$ roadside concentration between 2000 and 2017 in the FUAs in which the trends differ considerably from the overall European trends. The smoothed lines are loess (local regression) fits.

199 At the city resolution, while the majority of FUAs display a similar
 200 pattern in the NO_2 and NO_x concentration trends, the precise nature and
 201 location of the turning point in the NO_2 concentration trend varies, with
 202 the concentration in some FUAs increasing to a peak, before starting to
 203 decrease. It is possible that these differences in the shape of the trends are
 204 due to differences in the position of the NO_2 concentration peak resulting

205 from variations in the composition and age of the vehicle fleet between cities.

206 The exceptions to the pattern of monotonic decrease in NO_x concentration
 207 are Roma and Lisboa, where the NO_x concentration initially increases until
 208 around 2004-2005, then decreases. One possible explanation for this is a
 209 slower than average rate of vehicle turnover and an older vehicle fleet in
 210 Italy and Portugal compared with the rest of Europe (as shown in Table 2),
 211 resulting in a delay in the widespread presence of vehicles fitted with NO_x
 212 reducing emission technologies in these cities.

Table 2: The age of the vehicle fleet in Italy and Portugal, compared with the entire European Union vehicle fleet (ACEA, 2017).

Country	Vehicle type	Average age (years)	Proportion >10 years
EU	Passenger	10.7	0.49
EU	Light commercial	10.7	0.49
EU	Medium/heavy commercial	11.7	0.53
Italy	Passenger	10.7	0.52
Italy	Light commercial	11.9	0.59
Italy	Medium/heavy commercial	13.2	0.69
Portugal	Passenger	12.6	0.62
Portugal	Light commercial	14.0	0.79
Portugal	Medium/heavy commercial	14.0	0.79

213 Three FUAs possess trends in NO_2 concentration that differ considerably
 214 from the European-wide trend. As shown in Figure 6, in Madrid and Marseille
 215 NO_2 concentration increased to a peak in 2005-2006, then decreased before
 216 beginning to increase again around 2013. As with the FUAs exhibiting
 217 outlying NO_x trends, these trends can be interpreted as lagged versions of
 218 the common European trends due to variations in the composition of the
 219 vehicle fleet in these FUAs. However, the NO_2 concentration trend in Udine
 220 is distinct from that in any other FUA, with constant NO_2 concentration
 221 until 2007, decreasing to 2010, then increasing to a plateau in 2014. The
 222 reason for this is unclear.

223 As can be seen in Figure 5, the trends in PM_{10} concentration in all 24
 224 individual FUAs resemble the European-wide trend of monotonic decrease,
 225 with the single exception of Madrid, where PM_{10} concentrations increased
 226 to a maximum in 2005, before declining, as shown in Figure 6.

227 The initial increases in both NO_2 and PM_{10} concentration in Madrid
 228 could result from rapid growth in the number of diesel vehicles outstripping
 229 improvements in emission technologies. Between 1999 and 2008, the vehicle
 230 fleet in Madrid increased by 34%, with the proportion of the fleet composed

231 of diesel vehicles growing from 28% to 53% (Salvador et al., 2012). Compared
232 with the increase in the total number of vehicles in Europe of 14% between
233 2000 and 2008 (EEA, 2018), this represents a large increase in vehicles, and a
234 large increase in the number of diesel vehicles specifically. It is likely to have
235 led to an increase in vehicular NO_x and PM_{10} emissions which overwhelmed
236 the concurrent decreases in per-vehicle emissions.

237 Only six FUAs contained more than one monitoring site measuring $\text{PM}_{2.5}$
238 concentration that met the data capture requirements for the trend analysis.
239 In all six FUAs, $\text{PM}_{2.5}$ concentrations decreased over the period analysed.

240 4. Conclusions

241 A new trend evaluation approach has been applied to European roadside
242 ambient air pollution data that maximises data usage while minimising poten-
243 tially important biasing effects. Long term trends in roadside concentrations
244 of NO_x , NO_2 , PM_{10} and $\text{PM}_{2.5}$ in Europe and in 44 European functional
245 urban areas between 2000 and 2017 were evaluated using the rolling change
246 method. The city-scale trends for the most part displayed a strong consis-
247 tency with the overall European trends, namely monotonic decreases in NO_x
248 and $\text{PM}_{2.5}$ concentrations since 2000, and initial increases in NO_2 and PM_{10}
249 concentrations until 2002, followed by monotonic decreases. Reductions in
250 vehicle emissions resulting from improvements in vehicle exhaust technologies,
251 such as three-way catalysts, diesel particulate filters and diesel oxidation
252 catalysts are likely responsible for the observed declines in concentration.
253 In almost all cases where the city-scale trends differed considerably from
254 these patterns, the anomalous trends seemed to be lagged versions of the
255 same trend, possibly caused by variations in the traffic fleet composition. For
256 example, an older vehicle fleet, slower rate of vehicle turnover, or increases
257 in the number of vehicles (particularly diesel vehicles) may contribute to a
258 delay in the turning point of the trends in NO_x , NO_2 and PM_{10} roadside
259 concentration.

260 Acknowledgements

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

265 **References**

266 Carslaw, D. C., Farren, N. J., Vaughan, A. R., Drysdale, W. S., Young,
267 S., Lee, J. D., jan 2019. The diminishing importance of nitrogen dioxide
268 emissions from road vehicle exhaust. *Atmospheric Environment: X* 1,
269 100002.

270 EEA, 2016. Explaining Road Transport Emissions: A Non-Technical Guide.
271 Tech. rep., European Environment Agency.
272 URL [https://www.eea.europa.eu/publications/
273 explaining-road-transport-emissions](https://www.eea.europa.eu/publications/explaining-road-transport-emissions)

274 EEA, 2018a. Air Quality e-Reporting (AQ e-Reporting).
275 URL [https://www.eea.europa.eu/data-and-maps/data/
276 aqereporting-8](https://www.eea.europa.eu/data-and-maps/data/aqereporting-8)

277 EEA, 2018b. Air Quality in Europe. Tech. rep., European Environment
278 Agency.
279 URL [https://www.eea.europa.eu/publications/
280 air-quality-in-europe-2018](https://www.eea.europa.eu/publications/air-quality-in-europe-2018)

281 EEA, 2018. Size of the vehicle fleet.
282 URL [https://www.eea.europa.eu/data-and-maps/indicators/
283 size-of-the-vehicle-fleet/size-of-the-vehicle-fleet-9](https://www.eea.europa.eu/data-and-maps/indicators/size-of-the-vehicle-fleet/size-of-the-vehicle-fleet-9)

284 EEA, 2019. Eionet Central Data Repository. Tech. rep.
285 URL <http://cdr.eionet.europa.eu/>

286 Font, A., Fuller, G. W., nov 2016. Did policies to abate atmospheric emissions
287 from traffic have a positive effect in London? *Environmental Pollution*
288 218, 463–474.

289 Grange, S. K., 2016. smonitor: A Framework and a Collection of Functions
290 to Allow for Maintenance of Air Quality Monitoring Data.
291 URL <https://github.com/skgrange/smonitor>

292 Grange, S. K., Lewis, A. C., Moller, S. J., Carslaw, D. C., dec 2017. Lower
293 vehicular primary emissions of NO₂ in Europe than assumed in policy
294 projections. *Nature Geoscience* 10 (12), 914–918.

295 Lang, P., 2018. aqtrends.
296 URL <https://github.com/pollylang/aqtrends>

- 297 Lang, P. E., Carslaw, D. C., Moller, S. J., apr 2019. A trend analysis approach
298 for air quality network data. *Atmospheric Environment: X* 2, 100030.
- 299 Lavallo, C., Kompil, M., Aurambout, J.-P., 2015. UI - Boundaries for the
300 functional urban areas (LUISA Platform REF2014). Tech. rep., European
301 Commission, Joint Research Centre (JRC).
302 URL <http://data.europa.eu/89h/jrc-luisa-ui-boundaries-fua>
- 303 Masiol, M., Squizzato, S., Formenton, G., Harrison, R. M., Agostinelli, C., jan
304 2017. Air quality across a European hotspot: Spatial gradients, seasonality,
305 diurnal cycles and trends in the Veneto region, NE Italy. *Science of The
306 Total Environment* 576, 210–224.
- 307 Mavroidis, I., Chaloulakou, A., dec 2011. Long-term trends of primary and
308 secondary NO₂ production in the Athens area. Variation of the NO₂/NO_x
309 ratio. *Atmospheric Environment* 45 (38), 6872–6879.
- 310 Salvador, P., Artíñano, B., Viana, M., Alastuey, A., Querol, X., sep 2012.
311 Evaluation of the changes in the Madrid metropolitan area influencing
312 air quality: Analysis of 1999–2008 temporal trend of particulate matter.
313 *Atmospheric Environment* 57, 175–185.
- 314 The International Council on Clean Transportation, 2018. European vehicle
315 market statistics 2018/2019. Tech. rep.
316 URL <http://eupocketbook.theicct.org>

E. Appendix V: Supplementary material for Chapter 4

Table E.1: Metadata for the monitoring sites used in the analysis. The data was sourced from the AURN, LAQN and AQE networks.

Site code	Site name	Site type	Region	Pollutant
ABD	Aberdeen	Urban background	North East Scotland	no2, nox, o3
ABD7	Aberdeen Union Street Roadside	Urban traffic	North East Scotland	no2, nox
AH	Aston Hill	Rural background	North Wales	no2, nox
BAR3	Barnsley Gawber	Urban background	Yorkshire & Humberside	no2, nox
BATH	Bath Roadside	Urban traffic	South West	no2, nox
BEL1	Belfast Stockman's Lane	Urban traffic	Northern Ireland	no2, nox
BEL2	Belfast Centre	Urban background	Northern Ireland	no2, nox, o3
BRT3	Brighton Preston Park	Urban background	South East	no2, nox
BEX	London Bexley	Urban background	Greater London	no2, nox
BIL	Billingham	Urban background	North East	no2, nox
BORN	Bournemouth	Urban background	South West	no2, nox, o3
BR8	Bristol St Paul's	Urban background	South West	no2, nox, o3
BUSH	Bush Estate	Rural background	Central Scotland	no2, nox, o3
CA1	Camden Kerbside	Urban traffic	Greater London	no2, nox
CAM	Cambridge Roadside	Urban traffic	Eastern	no2, nox
CANT	Canterbury	Urban background	South East	no2, nox
CARD	Cardiff Centre	Urban background	South Wales	no2, nox, o3
CARL	Carlisle Roadside	Urban traffic	North West & Merseyside	no2, nox
CHP	Chepstow A48	Urban traffic	South Wales	no2, nox
CHS7	Chesterfield Roadside	Urban traffic	East Midlands	no2, nox
CLL2	London Bloomsbury	Urban background	Greater London	no2, nox, o3
CWMB	Cwmbran	Urban background	South Wales	no2, nox, o3
DUMF	Dumfries	Urban traffic	Scottish Borders	no2, nox
ECCL	Salford Eccles	Industrial	North West & Merseyside	no2, nox
HUL2	Hull Freetown	Urban background	Yorkshire & Humberside	no2, nox, o3

APPENDIX E. APPENDIX V: SUPPLEMENTARY MATERIAL FOR CHAPTER 4

ESK	Eskdalemuir	Rural background	Scottish Borders	no2, nox, o3
EX	Exeter Roadside	Urban traffic	South West	no2
FW	Fort William	Urban background	Highland	no2, nox, o3
HOPE	Stanford-le-Hope Roadside	Urban traffic	Eastern	no2, nox
HORE	Horley	Urban background	South East	no2, nox
GLA4	Glasgow Kerbside	Urban traffic	Central Scotland	no2, nox
GLAZ	Glazebury	Rural background	North West & Merseyside	no2, nox, o3
GRAN	Grangemouth	Industrial	Central Scotland	no2, nox
HG1	Haringey Roadside	Urban traffic	Greater London	no2, nox
HIL	London Hillingdon	Urban background	Greater London	no2, nox, o3
HM	High Muffles	Rural background	Yorkshire & Humberside	no2, nox, o3
HORS	London Westminster	Urban background	Greater London	no2, nox
HRL	London Harlington	Industrial	Greater London	no2, nox, o3
INV2	Inverness	Urban traffic	Highland	no2
KC1	London N. Kensington	Urban background	Greater London	no2, o3
LEAM	Leamington Spa	Urban background	West Midlands	no2, nox, o3
LED6	Leeds Headingley Kerbside	Urban traffic	Yorkshire & Humberside	no2, nox
LEED	Leeds Centre	Urban background	Yorkshire & Humberside	no2, nox, o3
LEOM	Leominster	Urban background	West Midlands	no2, nox, o3
LH	Lullington Heath	Rural background	South East	no2, nox, o3
REA1	Reading New Town	Urban background	South East	no2, nox, o3
LON6	London Eltham	Urban background	Greater London	no2, nox, o3
LVP	Liverpool Speke	Urban background	North West & Merseyside	no2, nox, o3
MAN3	Manchester Piccadilly	Urban background	North West & Merseyside	no2, nox, o3
MID	Middlesbrough	Urban background	North East	no2, nox, o3
MKTH	Market Harborough	Rural background	East Midlands	no2, nox
MY1	London Marylebone Road	Urban traffic	Greater London	no2, nox, o3
NEWC	Newcastle Centre	Urban background	North East	no2, nox, o3

APPENDIX E. APPENDIX V: SUPPLEMENTARY MATERIAL FOR CHAPTER 4

NOTT	Nottingham Centre	Urban background	East Midlands	no2, nox, o3
NPT3	Newport	Urban background	South Wales	no2, nox
OSY	St Osyth	Rural background	Eastern	no2, nox, o3
OX	Oxford Centre Roadside	Urban traffic	South East	no2, nox
OX8	Oxford St Ebbes	Urban background	South East	no2, nox
PEMB	Narberth	Rural background	South Wales	no2, nox, o3
PLYM	Plymouth Centre	Urban background	South West	no2, nox, o3
PMTH	Portsmouth	Urban background	South East	no2, nox, o3
PRES	Preston	Urban background	North West & Merseyside	no2, nox, o3
PT4	Port Talbot Margam	Urban background	South Wales	no2, nox, o3
SCN2	Scunthorpe Town	Industrial	Yorkshire & Humberside	no2, nox
SDY	Sandy Roadside	Urban traffic	Eastern	no2, nox
SHE	Sheffield Tinsley	Urban background	Yorkshire & Humberside	no2, nox
SOUT	Southampton Centre	Urban background	South East	no2, nox, o3
STOK	Stoke-on-Trent Centre	Urban background	West Midlands	no2, nox, o3
SUN2	Sunderland Silksworth	Urban background	North East	no2, nox, o3
SWA1	Swansea Roadside	Urban traffic	South Wales	no2, nox
TH2	Tower Hamlets Roadside	Urban traffic	Greater London	no2, nox
THUR	Thurrock	Urban background	Eastern	no2, nox, o3
WAR	Warrington	Industrial	North West & Merseyside	no2, nox
WFEN	Wicken Fen	Rural background	Eastern	no2, nox
WIG5	Wigan Centre	Urban background	North West & Merseyside	no2, nox, o3
WREX	Wrexham	Urban traffic	North Wales	no2, nox
YK11	York Fishergate	Urban traffic	Yorkshire & Humberside	no2, nox
YW	Yarner Wood	Rural background	South West	no2, nox, o3
CI1	Chichester - A27 Chichester Bypass	Urban traffic	South East	no2, nox
HB004	Watford Town Hall	Urban background	Greater London	no2, nox
HI3	London Hillingdon Oxford Avenue	Urban background	Greater London	no2, nox

APPENDIX E. APPENDIX V: SUPPLEMENTARY MATERIAL FOR CHAPTER 4

HO2	Horsham - Park Way	Urban traffic	South East	no2, nox
HT1	Hastings - Bulverhythe	Urban traffic	South East	no2, nox
KC2	Cromwell Road	Urban traffic	Greater London	no2, nox
KC3	Knightsbridge	Urban traffic	Greater London	no2, nox
KC4	Chelsea	Urban traffic	Greater London	no2, nox
KC5	Earls Court Road	Urban traffic	Greater London	no2, nox
MW1	Maidenhead Roadside	Urban traffic	South East	no2, nox
RG3	Reigate and Banstead Poles Lane	Rural background	South East	no2, nox, o3
T55	Heathrow Green Gates	Airport	Greater London	no2
ABD1	Aberdeen Anderson Dr	Urban traffic	North East Scotland	no2, nox
ABD3	Aberdeen Union Street Roadside	Urban traffic	North East Scotland	no2, nox
AYR	South Ayrshire Ayr High St	Urban traffic	Central Scotland	no2, nox
BRX	West Lothian Broxburn	Urban traffic	Central Scotland	no2, nox
CUPA	Fife Cupar	Urban traffic	Central Scotland	no2, nox
DUN5	Dundee Seagate	Urban traffic	North East Scotland	no2, nox
DUN6	Dundee Lochee Road	Urban traffic	North East Scotland	no2, nox
DUN7	Dundee Whitehall Street	Urban traffic	North East Scotland	no2
DUNF	Fife Dunfermline	Urban traffic	Central Scotland	no2, nox
EDB1	East Dunbartonshire Bishopbriggs	Urban traffic	Central Scotland	no2, nox
EDB2	East Dunbartonshire Bearsden	Urban traffic	Central Scotland	no2, nox
EDB3	East Dunbartonshire Kirkintilloch	Urban traffic	Central Scotland	no2, nox
FAL3	Falkirk Hope St	Urban traffic	Central Scotland	no2, nox
FALK	Falkirk Grangemouth MC	Urban background	Central Scotland	no2, nox
GLA6	Glasgow Byres Road	Urban traffic	Central Scotland	no2, nox
GLA7	Glasgow Waulkmillglen Reservoir	Rural background	Central Scotland	no2, nox, o3
NL3	N Lanarkshire Chapelhall	Urban traffic	Central Scotland	no2, nox
NL9	N Lanarkshire Moodiesburn	Urban traffic	Central Scotland	no2, nox
PET2	Perth Atholl Street	Urban traffic	North East Scotland	no2, nox

APPENDIX E. APPENDIX V: SUPPLEMENTARY MATERIAL FOR CHAPTER 4

PETH	Perth High Street	Urban traffic	North East Scotland	no2, nox
ROSY	Fife Rosyth	Urban traffic	Central Scotland	no2, nox
WDB3	West Dunbartonshire Clydebank	Urban traffic	Central Scotland	no2, nox
WDB4	West Dunbartonshire Glasgow Road	Urban traffic	Central Scotland	no2, nox
CAE4	Caerphilly White Street	Urban traffic	South Wales	no2, nox
CHEP	Chepstow A48	Unknown	South Wales	no2, nox
NPT1	Newport St Julians Comp School	Urban background	South Wales	no2, nox
SWA5	Swansea Morriston Roadside	Urban traffic	South Wales	no2, nox
CAS3	Castlereagh Dundonald	Urban traffic	Northern Ireland	no2, nox
CT6	City of London - Walbrook Wharf	Urban traffic	City of London	no2, nox
GR4	Greenwich - Eltham	Urban background	Greenwich	no2, nox, o3
TK1	Thurrock - London Road (Grays)	Urban background	Thurrock	no2, nox, o3
BG1	Barking and Dagenham - Rush Green	Urban background	Barking and Dagenham	no2
ZR1	Dartford Roadside - St Clements	Urban traffic	Dartford	no2, nox
ZR2	Dartford Roadside 2 - Town Centre	Urban traffic	Dartford	no2, nox
ZR3	Dartford Roadside 3 - Bean Interchange	Urban traffic	Dartford	no2, nox
ZV1	Sevenoaks - Greatness Park	Urban background	Sevenoaks	no2, nox, o3
ZY1	Canterbury Backgrnd - Chaucer TS	Urban background	Kent	no2
RG2	Reigate and Banstead - Horley South	Urban background	Reigate and Banstead	no2
TK3	Thurrock - Stanford-le-Hope	Urban traffic	Thurrock	no2, nox
BL0	Camden - Bloomsbury	Urban background	Camden	no2, nox, o3
LH0	Hillingdon - Harlington	Urban background	Hillingdon	no2, nox, o3
HI0	Hillingdon - Keats Way	Urban background	Hillingdon	no2, nox, o3
WM0	Westminster - Horseferry Road	Urban background	Westminster	no2, nox
CA2	Crawley - Gatwick Airport	Urban background	Crawley	no2, nox
GR9	Greenwich - Westthorne Avenue	Urban traffic	Greenwich	no2, nox, o3
NF5	New Forest - Lyndhurst	Urban traffic	New Forest	no2, nox
MW2	Windsor and Maidenhead - Clarence Road	Urban traffic	Windsor and Maidenhead	no2, nox

APPENDIX E. APPENDIX V: SUPPLEMENTARY MATERIAL FOR CHAPTER 4

BW1	Brentwood - Brentwood Town Hall	Urban background	Brentwood	no2, nox
BX1	Bexley - Slade Green	Urban background	Bexley	no2, nox, o3
BX2	Bexley - Belvedere	Urban background	Bexley	no2, nox
CD1	Camden - Swiss Cottage	Urban traffic	Camden	no2, nox
IS2	Islington - Holloway Road	Urban traffic	Islington	no2, nox
LW2	Lewisham - New Cross	Urban traffic	Lewisham	no2, nox
MD3	Central Beds - Sandy	Urban traffic	Herts & Beds	no2, nox
RG1	Reigate and Banstead - Horley	Urban background	Reigate and Banstead	no2, nox
RI1	Richmond Upon Thames - Castelnau	Urban traffic	Richmond	no2, nox
RI2	Richmond Upon Thames - Barnes Wetlands	Urban background	Richmond	no2, nox, o3
CR5	Croydon - Norbury	Urban traffic	Croydon	no2, nox
EN4	Enfield - Derby Road	Urban traffic	Enfield	no2, nox
GB6	Greenwich and Bexley - Falconwood	Urban traffic	Bexley	no2, nox, o3
GR7	Greenwich - Blackheath	Urban traffic	Greenwich	no2, nox
HK6	Hackney - Old Street	Urban traffic	Hackney	no2, o3
HR1	Harrow - Stanmore	Urban background	Harrow	no2, nox
HR2	Harrow - Pinner Road	Urban traffic	Harrow	no2, nox
EA6	Ealing - Hanger Lane Gytratory	Urban traffic	Ealing	no2, nox
ZV2	Sevenoaks - Bat and Ball	Urban traffic	Sevenoaks	no2, nox
TK8	Thurrock - London Road (Purfleet)	Urban traffic	Thurrock	no2, nox
GR8	Greenwich - Woolwich Flyover	Urban traffic	Greenwich	no2, nox, o3
ST6	Sutton - Worcester Park	Urban traffic	Sutton	no2, nox
RD0	Reading - New Town	Urban background	Reading	no2, o3
LL1	Wealden - Lullington Heath	Rural background	Wealden	no2, nox, o3
GN3	Greenwich - Plumstead High Street	Urban traffic	Greenwich	no2, nox, o3
PO1	Portsmouth Background AURN	Urban background	Hampshire	no2, nox, o3
SH1	Southampton - City Centre AURN	Urban background	Southampton	no2, o3
GN0	Greenwich - A206 Burrage Grove	Urban traffic	Greenwich	no2, nox