Investigating the Respiratory Rate Response to $PM_{2.5}$ Exposure in Asthmatic Adolescents

Sharan Maiya

4th Year Project Report
Computer Science and Mathematics
School of Informatics
University of Edinburgh

2020

Abstract

Estimating exposure-response relationships between air pollution exposure and health effects is important for driving policies for improving environmental health. India's capital of Delhi is one of the most polluted areas in the world and vulnerable groups such as asthmatic adolescents are particularly at risk of the detrimental effects of this pollution. Existing literature investigates associations between exposure to fine particulate matter $PM_{2.5}$ and poor health but does not reach conclusions of causality due to potential confounders and unsuitable data. In this project, we use data for asthmatic adolescents gathered in the DAPHNE study through the use of the AIRSpeck device to measure personal $PM_{2.5}$ exposure and the RESpeck device to measure respiratory rate and activity level in order to estimate exposure-response relationships. Using statistical methods of causal inference and causal discovery we demonstrate that $PM_{2.5}$ directly causes short-term (< 1 hour) changes in respiratory rate, adjusting for the confounders of activity level, temperature and relative humidity. This is the first study of its kind to demonstrate this direct dependence from observational data.

Acknowledgements

I would like to thank my supervisor Professor D.K. Arvind for his support and guidance throughout the project. I would also like to thank my colleagues and friends Zoë Petard, Teodora Georgescu and Nikita Nikolajev for their much valued advice during our weekly meetings.

Table of Contents

Chapter 1

Introduction

The World Health Organisation (WHO) launched the "Breathe Life" campaign in 2016 which asks cities across the planet to commit to achieving specific climate-related goals by the year 2030. One of these goals is to meet the WHO recommended guideline limit for average annual $PM_{2.5}$: $10\mu g/m^3$. $PM_{2.5}$ is known as fine particulate matter and refers to particles suspended in air with diameter less than $2.5\mu m$ across (roughly 3% the diameter of a human hair). These exist naturally in the world in dust however they are also a bi-product of sources of pollution such as vehicle exhausts, as are PM_{10} and PM_1 which are defined similarly. Research such as [42] and [41] has shown that long-term exposure to $PM_{2.5}$ increases mortality and hospital admissions - the WHO estimates in [35] that over 2 million people die every year due to complications caused by breathing indoor and outdoor air pollution.

In 2017 London became the world's first megacity to join the Breathe Life campaign and since then many more across the globe have followed suit. Delhi however, lags far behind. The city has infamously poor air quality with many parts of it reporting $PM_{2.5}$ levels in the "hazardous" range of $> 250\mu g/m^3$. Such levels are a serious threat to human health and Delhi is in desperate need of a push to implement aggressive strategies to overcome this threat. The ongoing "Delhi Air Pollution: Health aNd Effects" (DAPHNE) project aims to drive such policies by providing scientific evidence of the health effects of air pollution on vulnerable groups in the city such as pregnant mothers and asthmatic adolescents ([48]). The project is funded jointly by the UK Medical Research Council and the UK Natural Environment Research Council. A team of doctors and scientists from both the UK and India, led by the Centre for Speckled Computing at the University of Edinburgh, are working to estimate exposure-response relationships which describe the immediate physiological response to $PM_{2.5}$ exposure by monitoring both air quality and respiratory rate using novel sensors developed at the Centre. These sensors can comfortably be worn for observation periods and offer much more convenient methods than current equipment used in similar studies such as nasal cannulas ([34]).

It is clear that statistical evidence of causal relationships could add to the ever growing list of reasons to adopt better pollution-controlling policies. Knowledge of exposure-response relationships could also lead to better prediction of asthma exacerbations trig-

gered by exposure - this is important as young asthmatics are particularly susceptible to detrimental health effects associated with air pollution exposure as shown in studies such as [19].

1.1 Project Objectives

The aim of this project was to estimate exposure-response relationships between air pollution exposure and changes in respiratory rate in asthmatic adolescents in Delhi using the data gathered through the DAPHNE study. In this dissertation we employ two methods of causal inference in time series in order to do so: Granger Causality ([29]) and the PCMCI algorithm ([15]). Granger Causality is a concept adapted from the field of economics which asks whether one time series encodes unique and useful information for predicting another. PCMCI is a causal discovery method which learns causal relationships between variables in a multivariate system from data.

The DAPHNE study is in the middle of its data collection phase which began in 2018. These data are examined in their current state and the respiratory health effects of exposure to air pollution are explored using analytical methods. The following research questions are posed to be answered in this dissertation:

1. Can any insights be drawn regarding trends in air quality in Delhi and associations with respiratory rate?
2. Using these insights, what statistical analysis can be carried out in order to estimate exposure-response relationships from air pollution exposure to respiratory rate changes?
3. Do such methods prove or disprove the existence of a reliable short-term respiratory rate response in asthmatic adolescents to changes in $PM_{2.5}$ exposure i.e. a causal relationship?
4. If a relationship is found is it characterised by linear or non-linear dependencies?

1.2 Novel Contributions

This dissertation builds on previous work and introduces novel research by:

- being the *first* study of its kind to investigate a causal relationship between short-term ($<$ 1 hour) changes in $PM_{2.5}$ and respiratory rate at the *individual* level.
- focusing on *personal* exposure by monitoring air quality levels at the most geographically precise level possible.
- examining the *direct* relationship between $PM_{2.5}$ and respiratory rate by conditioning on confounders such as activity level, temperature and relative humidity.
- estimating exposure-response relationships and *how they evolve over time* by testing for effects of $PM_{2.5}$ changes at a series of different delays.

- carrying out an exhaustive search for *both linear and non-linear causal relationships* between $PM_{2.5}$ and respiratory rate using cutting-edge causal inference methods.

1.3 Report Structure

The structure of this dissertation is as follows:

- Chapter 2 reviews similar studies to DAPHNE as well as work carried out in the domain of estimating exposure-response relationships from observational data. We clarify how this project builds on previous work.
- Chapter 3 explains the background of the DAPHNE study and its aims as well as some technical basics of time series analysis used later in the report.
- Exploratory data analysis is carried out in **Part I** in order to understand trends in both exposure (Chapter 4) and respiratory rate (Chapter 5) observed for asthmatic adolescents in Delhi. This drives later decisions on data pre-processing and statistical analysis. The extent to which the method of data collection approximates a fully random study is also discussed. This is because conclusions of causality can only be drawn from unbiased results.
- A series of experiments are carried out in **Part II** to investigate exposure-response relationships in the data. Chapter 6 explains how the data are pre-processed before these experiments.
- Chapter 7 begins with simple analysis to motivate further investigation. We use partial dependence plots and regression analysis to look at associations between $PM_{2.5}$ and respiratory rate.
- Chapter 8 details the use of and results from tests for Granger Causality in the data. We use both parametric and non-parametric tests to look for all possible types of relationships and investigate the differences in results from different subjects/trials.
- Chapter 9 uses a powerful causal discovery method called PCMCI to estimate causal networks and model direct relationships between several variables measured in the data. It is paired with a linear and non-linear test for conditional independence in order to carry out an exhaustive search of all causal links in the data.
- Chapter 10 discusses the strengths and weaknesses of our approach and suggests future directions and improvements.

Chapter 2

Related Work

2.1 Health Effects of Air Pollution Exposure

Modelling the health effects of air pollution exposure using exposure-response relationships is essential to understanding its threat and driving policies to improve environmental health. In [23], the authors study the effects of air pollution on mortality rates in the US by regressing mortality rate on measures of total suspended particulate matter, sulphur dioxide and nitrogen dioxide while controlling for several other variables including age, sex and socio-economic status. The work found no significant link however it also concluded that regression analysis is inappropriate when assessing causal relationships - this contradicted many previous studies.

A decade later however the famous Harvard Six Cities Study ([13]) was published which linked air pollution to mortality in six US cities using Cox proportional-hazards regression modelling, a more complex non-linear model than the standard linear regressions used in [23]. The theme of contradicting air pollution epidemiology results from different studies has continued since then with conclusions usually depending heavily on the methodology of the analysis. In [42] Schwartz reaches the same conclusion as the Six Cities Study and establishes a causal link between air pollution and mortality using a Poisson regression model. He also uses a similar model to find an association between air pollution exposure and respiratory-related hospital admissions for the elderly in Detroit in [41].

In [28] Gamble argues that many conclusions reached in studies examining the effect of $PM_{2.5}$ exposure on mortality are spurious due to confounders such as activity level or lung function: this is an important point as it is often impossible to account for all potential confounders in an epidemiology study without significant domain knowledge. For example in [17] the authors conclude that the relationship between short-term exposure to $PM_{2.5}$ and mortality is not causal due to confounding by temperature. Examples like this are why randomisation across the study cohort is essential in order to reduce bias from confounders which may not be included in the model.

Machine learning methods have also aided researchers in carrying out these statistical analyses. In [11] the authors use target maximum likelihood estimation to establish a

causal relationship between air pollution exposure and low birth weight. The advantage of using a semi-parametric model like this is that fewer assumptions are made about the underlying system. Random forests and Bayesian networks are used to evaluate the relationship between short-term exposure to air pollution and cardiovascular-related hospital admissions among older adults in [5]. Other studies such as [24] and [26] report causal relationships between $PM_{2.5}$ exposure and increased risk of stillbirth and mortality. It should be noted however that the majority of these studies comment after their conclusions that further investigation is needed to confirm the exact exposure-response relationship being investigated. This is usually due to uncertainty in ruling out all possible confounders in the system either through a fully randomised study or through observation of them.

Epidemiology studies on observational data often analyse time series in order to draw conclusions about the delay between exposure and response. For this reason many techniques from economics research have been adapted and used. For example, both [26] and [43] investigate the long-term effects of $PM_{2.5}$ exposure on mortality using a difference-in-difference approach and Granger Causality respectively. Such methods are particularly applicable to the DAPHNE study as its data are large multivariate time series. We use Granger Causality in Chapter 8.

One of the earliest studies of the effects of air pollution exposure on the health of asthmatic patients is found in [16] where Schwartz finds a statistically significant association between short-term PM_{10} exposure and asthma-related hospital visits in Seattle. The main drawback of this study is of course its restriction to a relation of association from which no conclusions of causality can be inferred. The same is true of [20] where a positive (but not significant) association between PM_{10} exposure and frequency of asthma attacks is found. [20] improves on previous work by noting the problem of averaging citywide exposure indicators when measuring air pollution - this is not an accurate way of observing personal exposure. They instead use more geographically precise estimates. In [19] proximity to sources of air pollution is assessed for its effect on risk of asthma attacks in children. Among the study's conclusions is the finding that residence near a major source of air emissions is associated with an increased asthma attack risk of 108%.

We improve on these studies by

- using sophisticated methods of causal inference in order to assess whether the link between air pollution exposure and respiratory rate is causal or merely associative and/or driven by a confounder.
- carrying out multiple tests for causality for different delays in exposure in order to better estimate the time taken for the body to respond. This is possible due to the use of personal AIRSpeck and RESpeck monitors which are able to gather data at regular minute intervals - previous work tends to average over longer periods.
- measuring personal exposure at the most precise geographical level possible through the use of the personal AIRSpeck monitor which is worn by the subject during observation periods.

2.2 Deriving Exposure-Response Relationships

Exposure-response relationships describe the magnitude of the response of an organism as a function of exposure to a particular stimulus over time. Studying these within the context of drug doses, pollutants and other stimulants is important in order to establish environments which can be potentially dangerous to humans. Hence, assessing the health impacts due to any kind of environmental exposure is an important area of research.

This will inevitably involve making statements of causality and carrying out causal inference. In [7], [18] and [4] we find comparison and critique of various methods of causal inference particularly within the domain of epidemiology studies like DAPHNE. In [7] specifically Cox criticises the framework used by the United States Environmental Protection Agency (US EPA) for evaluating the existence of exposure-response relationships in studies as it is vague, relatively arbitrary and can easily be foiled with realistic counterexamples. He instead favours a statistical approach of empirical testing. Such tests can be evaluated using a variety of criteria including mutual information (which follows the intuition that causes are informative about their effects), directed dependence (which follows the intuition that information will flow from a cause to an effect) and internal / external consistency among others. A larger list of criteria is given in [6].

These criteria can be measured using many different methods. As mentioned in Section 2.1 regression analysis cannot be used to draw conclusions of causality however it can be used to understand potential directional associations: a useful initial step to take before causal inference. The field of economics has provided us with causal inference methods tailored to time series, the most well-known of these being Granger Causality. In addition, computationally intensive causal discovery methods learn a directed acyclic graph of causal relationships from data. These methods are very new: current literature is generally focused either on laying the theoretical groundwork for their use or proof of correctness. For example [45] summarises the broad concepts underlying these methods and the direction research is moving in while [25] demonstrates that causal discovery methods can learn the ”gold standard” of Alzheimer’s Disease knowledge from observational data. We employ all three of these techniques in **Part II**.

Chapter 3

Background

3.1 The DAPHNE Study

The DAPHNE study is a consortium of doctors and scientists from institutions both in the United Kingdom and India to study the effects of air pollution in Delhi on the health of pregnant mothers and their newborn children (MC cohort) as well as asthmatic adolescents (AP cohort). This dissertation is limited to the latter group.

This is done by monitoring both the personal exposure and respiratory rate of participants using sensors developed at the Centre for Speckled Computing at the University of Edinburgh.

3.1.1 The AIRSpeck Sensor

Air pollution data are gathered through the AIRSpeck sensor which comes in two variants: stationary and personal. These are shown in Figures 3.1 and 3.2. Both use an optical particle counter to monitor concentrations of PM_1, $PM_{2.5}$ and PM_{10} as well as temperature and relative humidity. The stationary sensor reports observations at 5 minute intervals while the personal sensor does so every minute. These are time stamped and tagged with GPS location, resulting in a spatiotemporal time series of regular air quality measurements ([2]).

Figure 3.1: The stationary AIRSpeck sensor uses a solar panel for power and can be affixed to common street furniture such as a pole or lamp post as shown.

(a) The sensor placed on a table.

(b) The sensor is worn on the belt as shown during observation periods.

Figure 3.2: The Personal AIRSpeck monitor used to measure exposure to polluted air in the DAPHNE study.

3.1.2 The RESpeck Sensor

Respiratory rate is monitored with the RESpeck device. Measuring 4.5 x 3.7 x 1.3cm and weighing only 17 grams this small lightweight sensor contains a tri-axial ac-

celerometer which can be used to calculate a respiratory rate signal. The method of calculation and validation against a nasal cannula are provided in [3]. The sensor can be worn on the skin and secured with medical tape as shown in Figure 3.3. When worn as shown below the ribs on the left side of the chest it can measure both respiratory rate as well as activity level due to close proximity to the centre of mass of the human body, as explained in [14].

(a) The RESpeck device placed beside a 50p coin.

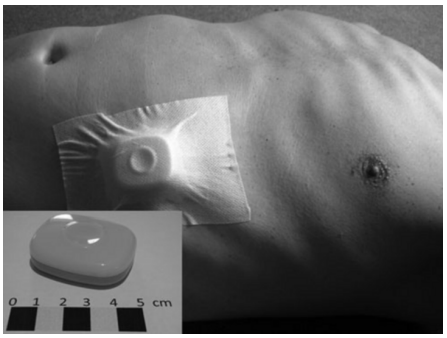

(b) The RESpeck device is worn on the skin and secured with medical tape for respiratory rate and activity level monitoring.

Figure 3.3: The RESpeck device used to measure respiratory rate in the DAPHNE study.

The combined use of personal AIRSpeck and RESpeck sensors allows for continuous monitoring of air pollution intake and the resulting effect on respiratory rate ([1]).

3.1.3 Aims

The principle aim of the DAPHNE study is to estimate exposure-response relationships of exposure to air pollution and health effects, particularly in vulnerable groups such as asthmatic adolescents. Gaining a better understanding of these relationships can help provide advice both on a personal level and policy-wide level on how best to mitigate the detrimental effects of exposure.

The DAPHNE project began in 2016 and is funded to continue until 2021. Over this 5 year period, 240 asthmatic adolescents are expected to take part in the study; each will complete three 48 hours observation periods (trials) during which their exposure to air pollution and respiratory rate are observed using the AIRSpeck and RESpeck sensors. Trials from the same subject are scheduled to take place during different seasons. At the time of writing, data from 125 subjects (183 individual trials) has been collected for the AP cohort; the earliest trial was conducted in September 2018 and the latest in February 2020. Hereafter trials are referred to by their ID which follows the format: DAP[subject number]([trial number]) e.g. DAP001(1) is subject 001's first trial.

Details of data collection are found in **Part I**.

3.2 Time Series Analysis

In **Part II** the DAPHNE data are analysed using statistical methods of causal inference developed for use with time series data. We now define some of the standard terminology used.

3.2.1 Definitions

- **Time Series**: A time series $X = (X_1, \ldots, X_n)$ is a sequence of real values taken at successive equally spaced intervals in time from $t = 1$ to $t = n$.
- **Lag**: Given a time series X and an observation x_i, the observation p intervals back x_{i-p} is called the p-th lag of x_i.
- **Stochastic Process**: A stochastic process is a family of indexed random variables defined on the same probability space. It is common to consider a time series a realisation of a stochastic process.
- **White Noise**: A white noise process is an uncorrelated stochastic process with zero mean and finite variance. If each variable in the process follows a normal distribution we can refer to it as Gaussian white noise.

3.2.2 Stationarity

Stationarity is an important property which deserves careful explanation as it is often a requirement for time series analysis. Indeed the Granger Causality tests used in Chapter 8 require stationary data and this must be tested for.

Stationarity is essentially the idea that the statistical properties of a process generating a time series do not change over time. Formally, a stochastic process X is said to be stationary if for any finite sub-sequence $\{X_s : s \in S\}$ the joint distributions of X_S and a time-shifted sequence X_{S+p} are the same. In practice this is impossible to verify however a weaker similar property is often all that is needed for statistical tests. This is referred to as *weak* stationarity; it requires that the process have constant mean and that the covariance of any two observations only depends on the absolute distance (in time) between them. Hereafter stationarity will mean weak stationarity when used without comment [1]. Figure 3.4, taken from [55], illustrates some of the clear differences between plots of stationary and non-stationary processes.

It is common to pre-process a time series by removing trend and seasonality in order to obtain stationary data. We do so in Chapter 6. Establishing whether it is necessary to do so requires a test for stationarity - many of which have been developed. These tests offer a more rigorous approach than just inspecting a plot of the series. Most are parametric: certain assumptions are made about the data generating stochastic process which allow us to test for properties which should hold should the process be stationary. Many have also been implemented in popular data-processing languages such as Python or R.

[1] It should be noted that the original definition of stationarity does not imply weak stationarity, nor does the converse hold.

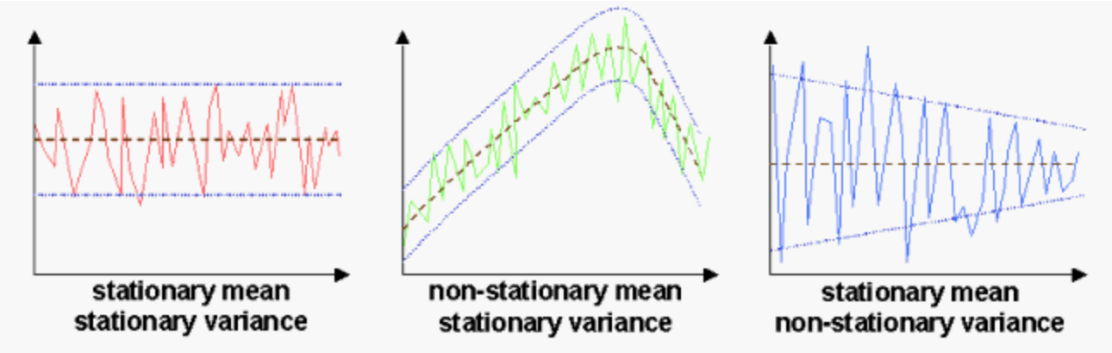

Figure 3.4: Stationary and non-stationary processes

In Section 8.1 we carry out VAR-based tests for Granger Causality which require us to test for (and obtain if necessary) stationary data. The Augmented Dickey-Fuller test ([56]) was developed for this exact purpose: it assumes an autoregressive model and tests for a unit root - the existence of which would imply non-stationarity. This test has a convenient implementation in the Python *statsmodels* package ([54]).

Section 8.2 continues by exploring tests for non-linear Granger Causality. Ideally an appropriate non-parametric test for stationarity should be paired with such analysis however the reality is more complicated. To our knowledge, no applicable test exists as current research has only achieved tests for specific types of data or processes. For example [32] proposes a test applicable to univariate time-homogenous Markov processes and [21] suggests a test for zero-mean discrete time random processes based on local Fourier analysis. None of these tests offer implementations either, hence (unfortunately) non-linear stationarity is assumed without proof in this section.

Part I

Exploratory Data Analysis

Chapter 4

The AIRSpeck Data

Air pollution data for the DAPHNE study is collected through the use of stationary and personal AIRSpeck sensors: these obtain geographically precise measurements of exposure at minute intervals. Figure 4.1 plots the $PM_{2.5}$ concentration for trial DAP001(2) as measured by their personal AIRSpeck monitor.

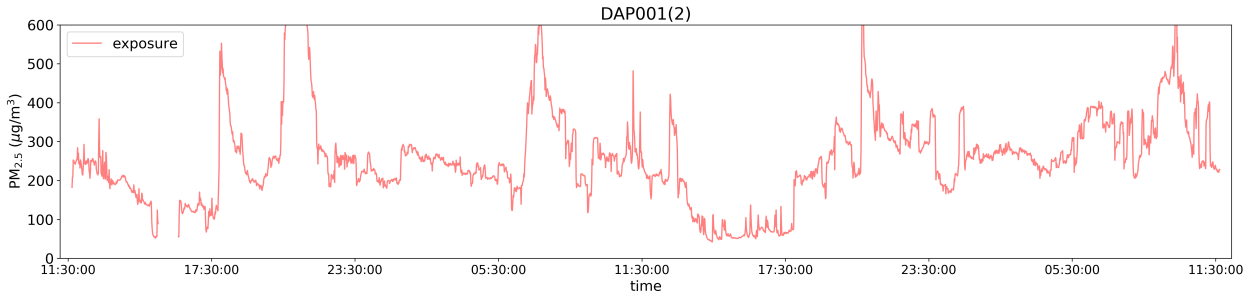

Figure 4.1: An example of $PM_{2.5}$ levels observed for one trial using the personal AIRSpeck.

The AIRSpeck sensor gathers information on the local air quality. Table 4.1 lists the different features of its reported data.

Table 4.1: Explanation of features measured by AIRSpeck sensors.

Feature	**Explanation**
Timestamp	UTC timestamp of current observation.
PM_1	Particulate matter with diameter $< 1\mu m$ ($\mu g/m^3$).
$PM_{2.5}$	Particulate matter with diameter $< 2.5\mu m$ ($\mu g/m^3$).
PM_{10}	Particulate matter with diameter $< 10\mu m$ ($\mu g/m^3$).
Temperature	Temperature of surrounding air ($^\circ C$).
Relative Humidity	% moisture in air after adjusting for temperature.
Latitude	GPS latitude coordinates.
Longitude	GPS longitude coordinates.

Due to GPS location tagging on personal AIRSpecks, it is possible to paint a picture of the air quality at different locations in Delhi at different points in time as shown in Figure 4.2. This helps to provide an accurate measurement of personal exposure. In addition, at least two stationary AIRSpeck sensors are set up to gather more information during each trial: one inside the subject's house and another set up outdoors near their school. In some cases a third stationary AIRSpeck is set up in the subject's local community. Figure 4.3 shows the various locations of these static monitors. It is important that the subjects live in different communities across Delhi - this increased randomisation helps to avoid spurious results which can arise due to specific trends in exposure in certain parts of the city. We ensure the same distribution across all subjects and account for local variations in personal exposure to air pollution by standardising all data. This is necessary as some subjects may be more used to higher levels of $PM_{2.5}$ exposure than others - we wish to examine the health effects of exposure at the individual level so we must standardise to make results between different subjects comparable. This is explained further in Chapter 6.

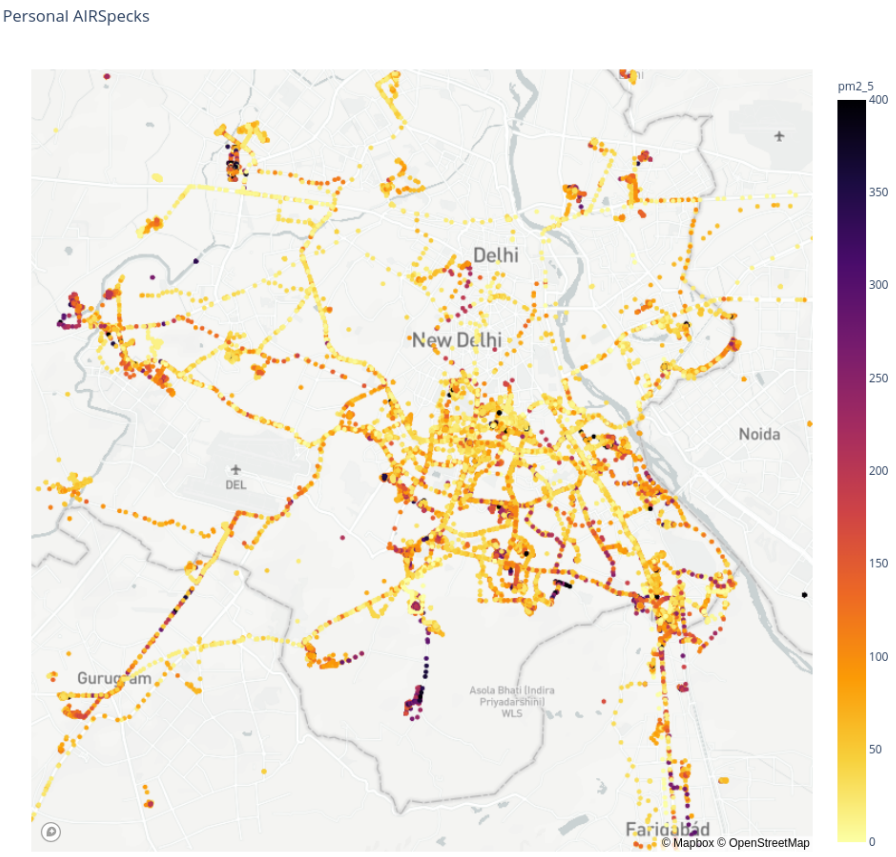

Figure 4.2: All personal AIRSpeck data collected at the time of this project. Observations are colour-coded to illustrate the variance in air pollution exposure, with darker points signifying higher levels of $PM_{2.5}$.

Part II examines the relationship between changes in $PM_{2.5}$ measured by the personal

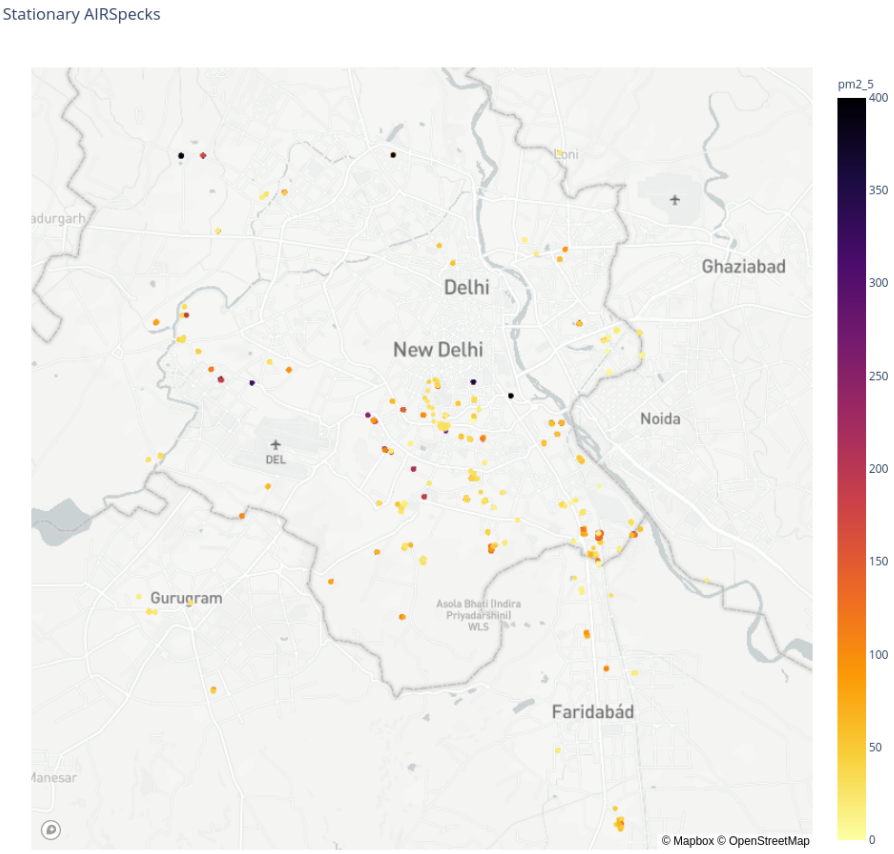

Figure 4.3: All stationary AIRSpeck data collected at the time of this project. This includes indoor sensors in the home and outdoor sensors at schools and in local communities. Darker observations signify higher levels of exposure.

AIRSpeck monitors and respiratory rate, in order to focus on personal exposure.

The distribution of $PM_{2.5}$ levels measured by personal AIRSpecks across all 183 trials is shown in Figure 4.4. The plot peaks close to the lower end of its range: high spikes in $PM_{2.5}$ are generally rare as they occur only in specific situations e.g. proximity to a source of great pollution such as a large industrial factory. A WHO study conducted in 2013 reported the average $PM_{2.5}$ levels in Delhi to be $153\mu g/m^3$. This is higher than the mean of the DAPHNE data which sits around $100\mu g/m^3$, with the peak of the distribution at an even lower value of roughly $45\mu g/m^3$. This discrepancy can be explained by considering the nature of the data collected. All subjects are adolescents and will therefore spend most of their day in school, returning home in the evening. These indoor environments will shield them from the polluted outdoor air reducing their overall exposure. Additionally, a Centre for Science and Environment (CSE) study recently found that yearly average pollution levels in Delhi have been decreasing, indicating that the WHO-reported average of $153\mu g/m^3$ is likely outdated. Nonetheless this distribution highlights that Delhi is still an extremely polluted city; recall that the WHO guideline for healthy annual $PM_{2.5}$ levels is $10\mu g/m^3$.

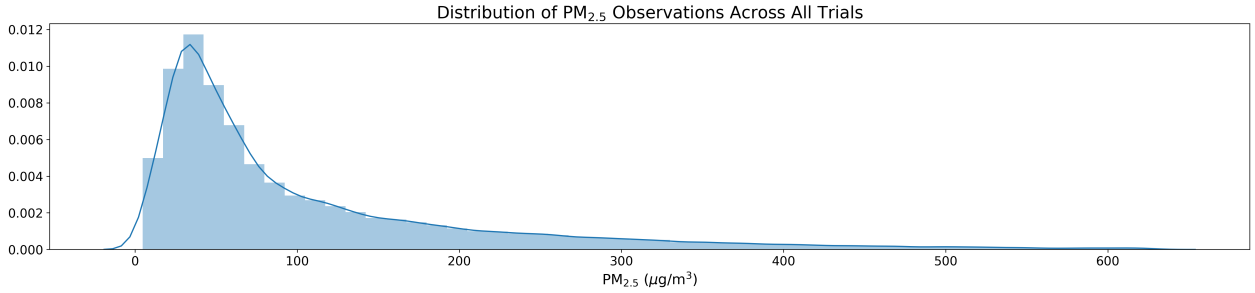

Figure 4.4: Distribution of $PM_{2.5}$ observations measured by personal AIRSpeck sensors across all 183 trials. The plot includes a Gaussian kernel density estimate (line).

By the end of the DAPHNE study each subject will have taken part in a total of three trials which must take place during different seasons of the year. This is done to mitigate the effects of seasonal trends in exposure which are very apparent in the data as shown in Figure 4.5. This plots the varying distribution of $PM_{2.5}$ (recorded on personal AIRSpeck sensors) over the months of the calendar year. For Delhi, this can be divided into five seasons. The months of December and January are considered winter. The period from February to March is spring. Summer continues from April to June. The monsoon season follows from July to September. Finally we consider the months of October to November as the ”burning” season as this roughly coincides with the period during which farmers burn their stubble after harvest, which undoubtedly has an effect on local air quality.

Figure 4.5 shows that the air quality is generally poorer during the colder months from November to February: the average $PM_{2.5}$ level is at its highest and distribution is at its flattest. This is partly due to the burning of biomass during the burning season however it is further driven by a meteorological phenomenon known as an inversion layer. Consider the rising of warm air in summer - this naturally carries away pollutants from the Earth’s surface. During colder months (and at night on a smaller timescale) cooler air closer to the ground becomes trapped underneath a layer of warmer air above which acts as a lid. Pollutants produced also become trapped and continue to mix with the air. The problem is exaggerated due to increased emissions in colder months for reasons such as heating. These strong seasonal trends in the data have the potential to bias statistical results should all trials be carried out during the same time of year.

As mentioned, inversion layers can also be observed at night. Figure 4.6 shows the change in distribution of $PM_{2.5}$ levels over the course of the 24 hours of the day as measured only by *outdoor* stationary AIRSpecks. $PM_{2.5}$ values are noticeably lower and less varied during the afternoon i.e. the warmest part of the day. During the night an inversion layer traps cold air along with airborne particulates close to the Earth’s surface[1].

The effects of the nightly decrease in air quality are mitigated to an extent by staying indoors. Figure 4.7 plots the change in $PM_{2.5}$ per hour as measured by indoor station-

[1]This is why citizens of large polluted cities are often advised against early morning walks or exercise as this is when their potential exposure to air pollution is at its highest.

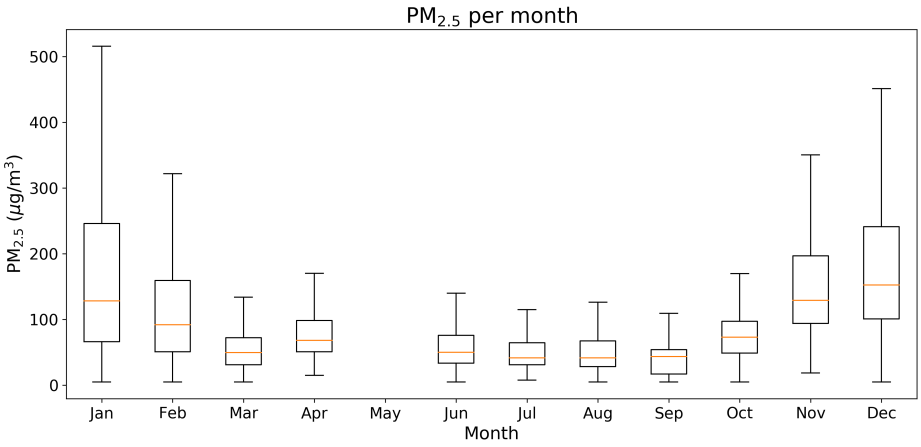

Figure 4.5: The changing distribution of $PM_{2.5}$ over the calendar year as measured by personal AIRSpecks. There is a clear trend of higher exposure and variability during colder months. It should be noted that no trials have been completed during the month of May.

ary AIRSpecks. Here a peak both in trend and variability is observed around 9am - this coincides roughly with the time when the house may have more particulate matter due to cooking of meals. It may also be the case that traffic-related pollution increases and enters the home through open windows or doors. A second smaller peak is observed around 12 hours later which again could be caused by analogous reasons or particulate matter generated due to the heating of the house as it gets colder towards the night. Levels are generally lower during the daytime when it is warmer. These plots highlight that personal exposure to $PM_{2.5}$ can vary considerably over the course of the day depending on the environment the subject is in. For this reason each trial in the DAPHNE study is conducted over a 48 hour period in order to provide more comparable results across different subjects through increased randomisation of exposure.

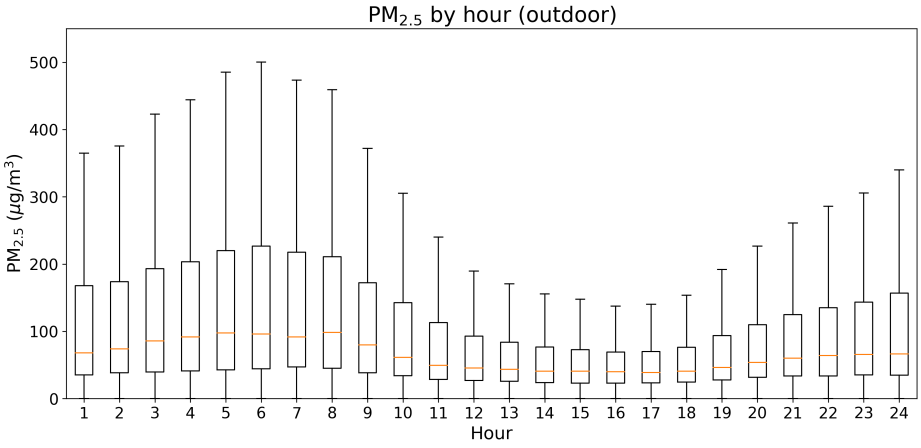

Figure 4.6: The changing distribution of $PM_{2.5}$ over the 24 hours of the day as recorded by outdoor stationary AIRSpecks.

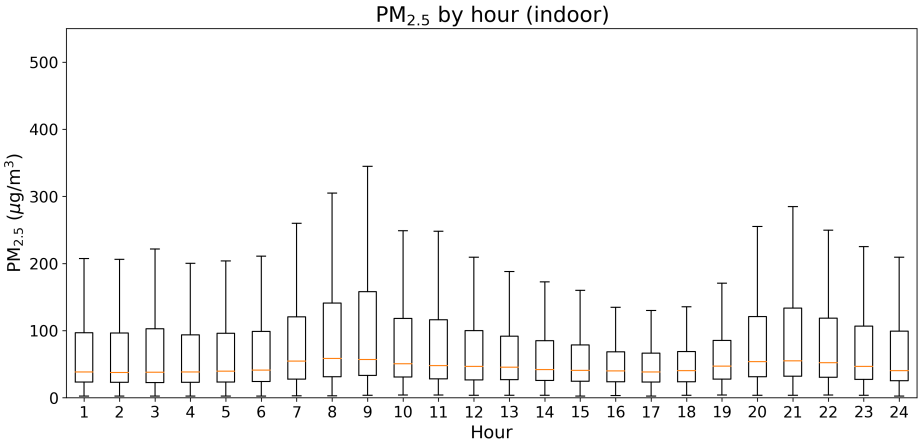

Figure 4.7: The changing distribution of $PM_{2.5}$ over the 24 hours of the day as recorded by indoor stationary AIRSpecks.

Chapter 5

The RESpeck Data

The respiratory rate of asthmatic adolescents is observed in the DAPHNE study using the RESpeck sensor, also at minute intervals. Figure 5.1 shows the respiratory rate observations for trial DAP088(1) over the full period.

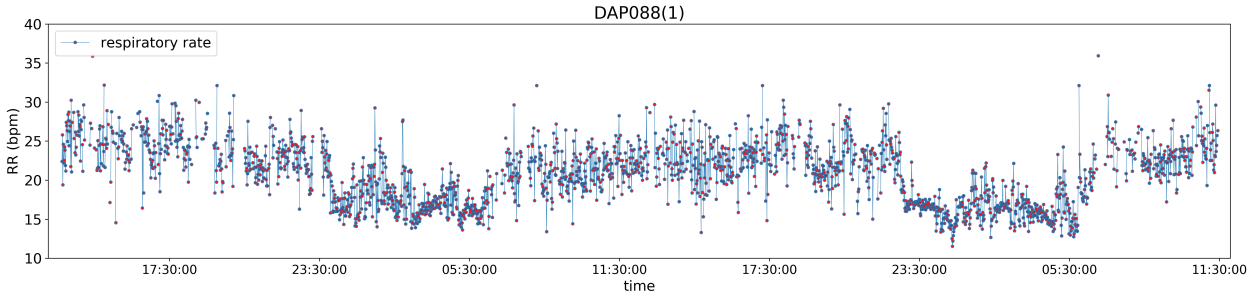

Figure 5.1: An example of respiratory rate observations for one trial.

Respiratory rate is not the only quantity observed by the RESpeck sensor. Table 5.1 lists the different features measured through the use of this device[1].

Eupnea is defined informally as ”normal” breathing or breathing at rest. While the exact figures vary depending on the source it is generally agreed that eupnea for a teenager or young adult human should be around 12-20 breaths per minute. However Figure 5.2, which plots the distribution of observed respiratory rate values across all

[1]For an explanation of activity level measurement see [14].

Table 5.1: Explanation of features measured by RESpeck sensors.

Feature	**Explanation**
Timestamp	UTC timestamp of current observation.
Respiratory Rate	Minute-averaged breathing rate.
Respiratory Rate: ST.D.	Standard deviation of breathing rate.
Activity Level	Measurement of activity on a continuous scale.
Activity Type	Discrete activity types inc. sitting/lying/walking/etc.

trials from all subjects, disagrees with this range - eupnea for this cohort seems to be between 15 and 25 breaths per minute. This higher average respiratory rate cannot be attributed to asthma: while research such as [22] indicates that respiratory rate does increase in naturally occurring acute asthma the consensus is that resting breathing rate should still sit within the range of 12-20 bpm ([49]). The other commonality between all subjects is their location in Delhi, one of the most polluted cities in the world. It is possible that living in such conditions has caused a natural increase in their resting respiratory rate however to our knowledge no research has been carried out to investigate this.

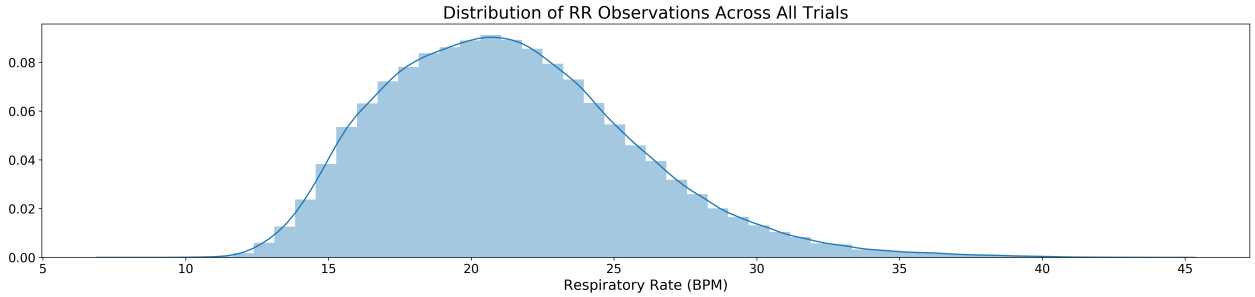

Figure 5.2: Distribution of respiratory rate observations measured on RESpeck sensors across all 183 trials. The plot includes a Gaussian kernel density estimate (line).

Figure 5.3 plots the distribution of recorded respiratory rate for a random subset of 50 subjects in the cohort. It shows clear and significant differences between subjects which are likely driven by a variety of factors such as age, sex, weight, lung function and general health. Aside from the commonality in adolescence and asthma, the cohort is randomised across all these dimensions. However this analysis indicates that "normal" breathing rate differs at the subject-level and this dictates the level at which any statistical analysis must take place i.e. exposure-response relationships (should they exist) should be estimated for each subject or possibly even each individual trial.

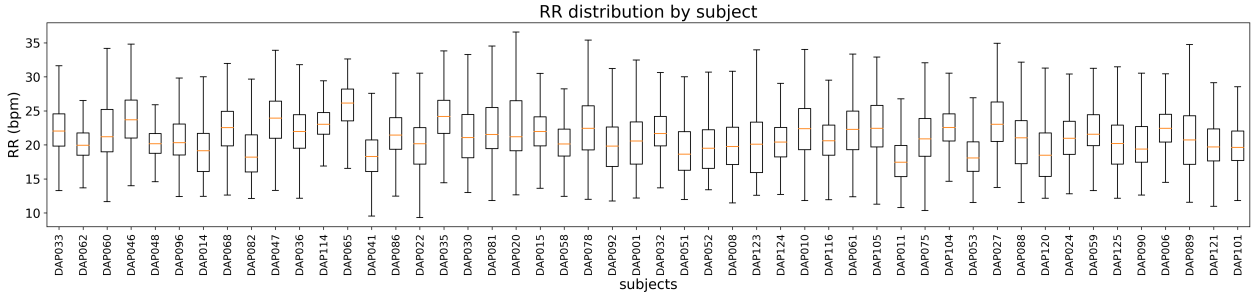

Figure 5.3: Distribution of respiratory rate observations for random subset of subjects (N=50).

Working at the subject/trial level means working with multiple smaller datasets as opposed to one large dataset. When using methods which learn from or fit to data the sample size must be large enough to allow the method to generalise well to new unseen data. Figure 5.4 shows the total number of minutes of respiratory rate observations for

the same subset of subjects. While some variations in the plot are due to the asymmetry in number of trials completed, it is also the case that many timestamps are missing. This problem, its potential consequences and a partial solution are discussed in Chapter 6.

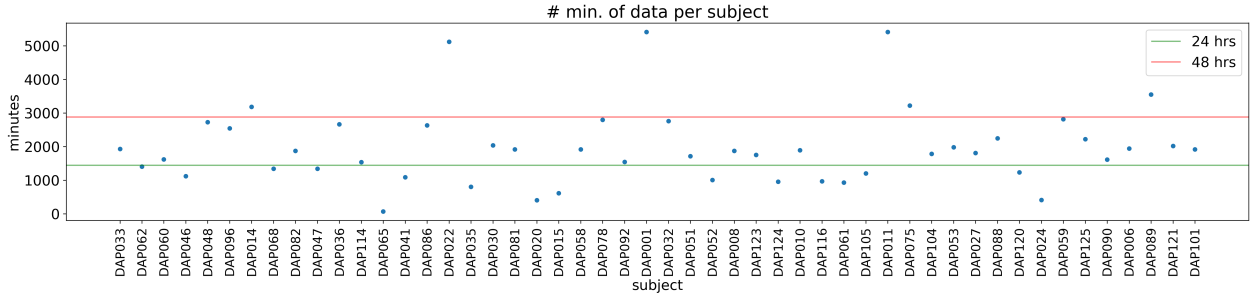

Figure 5.4: Number of minutes of observed respiratory rate values for the same subset of subjects shown in Figure 5.3.

Activity level is suggested as a potential confounder in air pollution epidemiology studies in [28]. The RESpeck sensor can be used to measure this and also classify between different types of activity e.g. lying/sitting/standing. Figure 5.5 plots both respiratory rate and activity level for trial DAP026(1) - we see a clear association between the two which is also seen in all other trials. The periods of lowest activity are during the night when the subject is lying down and sleeping - respiratory rate during these hours is also clearly lower both in trend and volatility. During the day when the subject is awake and more active their respiratory rate becomes more varied and generally higher. In Chapter 4 it is noted that the trend of $PM_{2.5}$ increases at night due to an inversion layer in the atmosphere. Therefore, since activity level is so closely associated with time of day, it has a common association with both $PM_{2.5}$ and respiratory rate and must be accounted for when carrying out statistical analysis. This is done via local standardisation, further explained in Chapter 6.

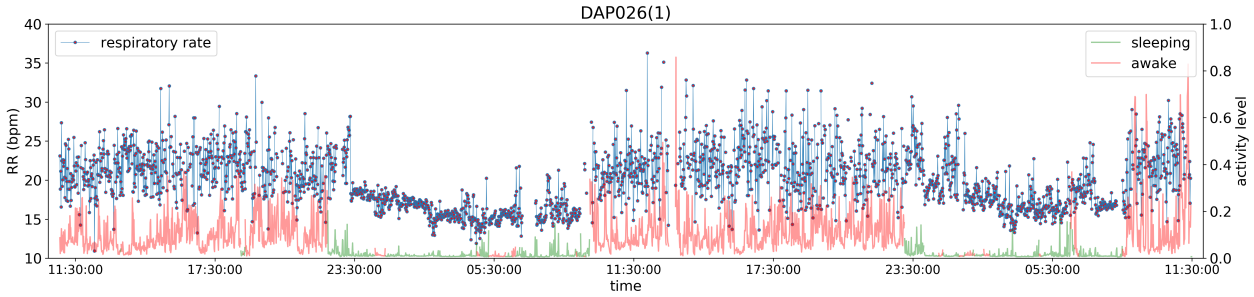

Figure 5.5: Respiratory rate (blue) and activity level (red/green) over the course of one trial. Green activity level corresponds to periods during which the classified activity type was ”lying down”; red is otherwise. An association in trend between activity level/type and respiratory rate is observed.

Part II

Experiments and Results

Chapter 6

Data Pre-Processing

This dissertation aims to thoroughly examine the DAPHNE data for the existence of a causal relationship between $PM_{2.5}$ and respiratory rate. Partial dependence plots and regression analysis motivate further investigation. Then techniques of causal inference are employed: in particular the concept of Granger Causality which asks whether one time series is useful in forecasting another. Standard VAR-based models and parametric tests investigate a possible linear relationship between the variables, while non-parametric methods which impose fewer restrictions can uncover more complex relationships. Both types of relationship may exist, hence both types of test should be used. Finally cutting edge causal discovery methods are used to learn causal networks from the data as these show much promise when it comes to identifying direct causal relationships and controlling for false-positives. Statistical tests find significant causal relationships ($p < 0.05$) between lags of $PM_{2.5}$ and respiratory rate.

Before carrying out such analysis however, the data must be cleaned and pre-processed.

6.1 Calibration

The use of multiple AIRSpeck sensors to collect data on PM_1, $PM_{2.5}$ and PM_{10} data can be problematic. Experimental error due to small differences between sensors is controlled for by calibrating all sensors against a reference. This is done by testing them before the study in a controlled environment such as a laboratory and comparing the recorded data. Calibration factors can then be determined for each sensor in order to calibrate its data by linear scaling - previous work at the Centre for Speckled Computing at the University of Edinburgh has shown this can be done as long as the relative humidity is below 80%. If this is not the case a more complex calibration method is required - this is achieved through the use of a trained neural network. In order to ensure accurate results in this project, AIRSpeck sensor data are calibrated before any further cleaning, pre-processing or analysis is completed.

6.2 Anomalous Observations

Anomaly detection in observational data is a difficult task: determining whether an extreme value is an error or just a surprise is important as there is no need to correct for values which are already correct. However there are some clearly unrealistic observations in the data - an example is shown in Figure 6.1 which plots personal exposure to $PM_{2.5}$ for trial DAP120(1) and shows a massive spike observed towards the end of the trial. Ambient $PM_{2.5}$ exposure does not reach heights of $28000\mu g/m^3$. When the spike is ignored, $PM_{2.5}$ values fall within a more realistic range as shown in Figure 6.2.

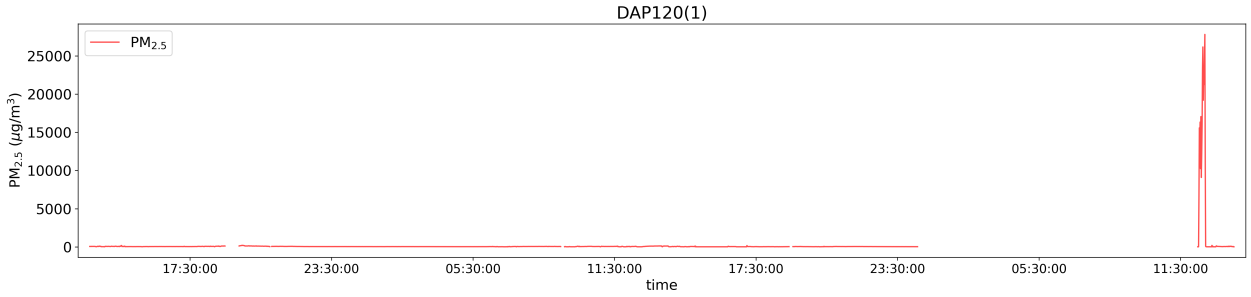

Figure 6.1: $PM_{2.5}$ observations with a massive spike

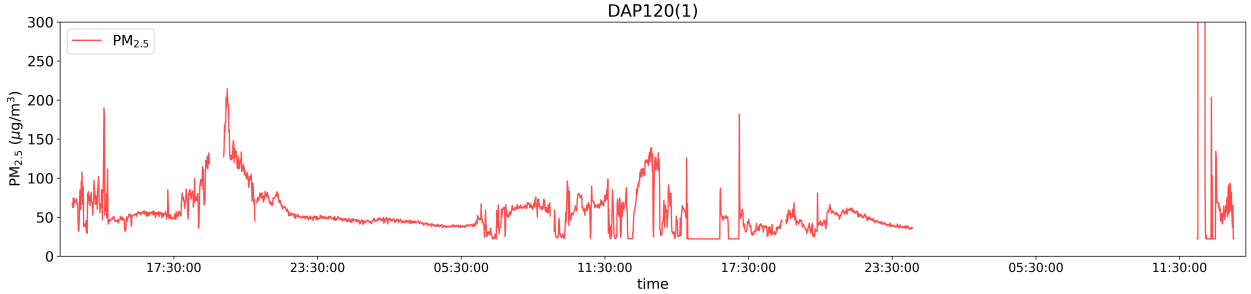

Figure 6.2: $PM_{2.5}$ observations ignoring the massive spike

A common technique used when dealing with observational data in statistics is called *winsorizing*. This is a method of anomaly removal by limiting or clipping extreme values in the data. For the DAPHNE data, a 90% winsorization is carried out: all values below the 5th percentile are set to the 5th percentile and similarly all values above the 95th percentile are set to the 95th percentile for features which have clear anomalies like $PM_{2.5}$. This is done on a per trial basis due to the large differences observed in data distributions between subjects noted in Chapter 5.

6.3 Interpolation

As is often the case when it comes to observational data, many trials are incomplete with missing values. Causes include subjects removing their wearable RESpeck at night as seen in Figure 6.3, a sensor running out of battery power or even a sensor malfunction. In addition, RESpeck recordings during periods of high activity level are

dropped due to high inaccuracy in respiratory rate calculation. These problems are inevitable but the handling of them is a crucially important task. We opt to interpolate as much data as possible without significant loss of information.

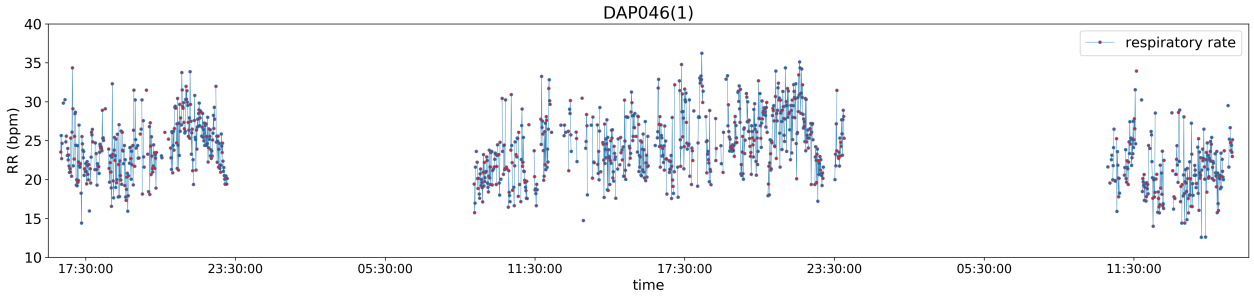

Figure 6.3: Respiratory rate observations in this trial clearly indicate the subject removed their RESpeck sensor during the night.

Interpolation must be carried out very carefully in order to avoid creating artificial unrealistic results. Intuitively, interpolating power decreases the further into the future we extrapolate i.e. for data observed at minute intervals, a gap of five minutes can be imputed with greater accuracy than a gap of five hours. Hence it is logical to set a threshold of time above which it no longer makes sense to impute missing values due to large error. This threshold is tuned for both RESpeck and AIRSpeck data using Algorithm 1.

Algorithm 1: Interpolation Threshold Tuning

input : max_interpolate = maximum gap of missing data over which to interpolate
output: Mean absolute percentage error of interpolation at this threshold

num_iterations $>> 1$;
good_trials = subset of trials with minimal missing data (e.g. $< 10\%$ missing data);
data = data belonging to good_trials;
for $i \leftarrow 1$ **to** num_iterations **do**
    randomly remove chunks of data up to size max_interpolate;
    interpolate all gaps up to size max_interpolate;
    calculate mean absolute percentage error from original data;
calculate average mean absolute percentage error;
return average mean absolute percentage error;

Figure 6.4 plots the results of running Algorithm 1 using a linear interpolation method for respiratory rate values recorded on RESpeck devices for thresholds from 1 to 60 minutes, while Figure 6.5 plots the same for AIRSpeck $PM_{2.5}$ observations. A reference line is added to both plots at a percentage error of 0.1%, which is chosen as the cutoff. The plots determine that the largest gaps over which RESpeck and AIRSpeck data can be interpolated accurately are 8 minutes and 14 minutes respectively. This ex-

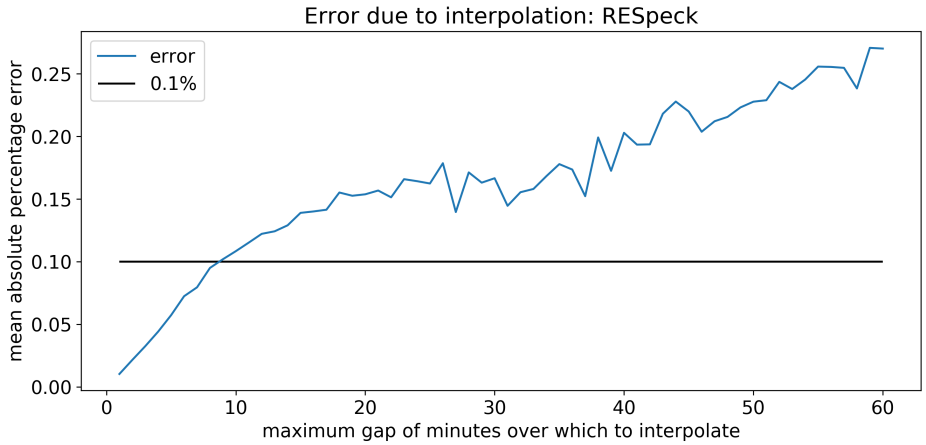

Figure 6.4: Percentage error of interpolation of RESpeck values

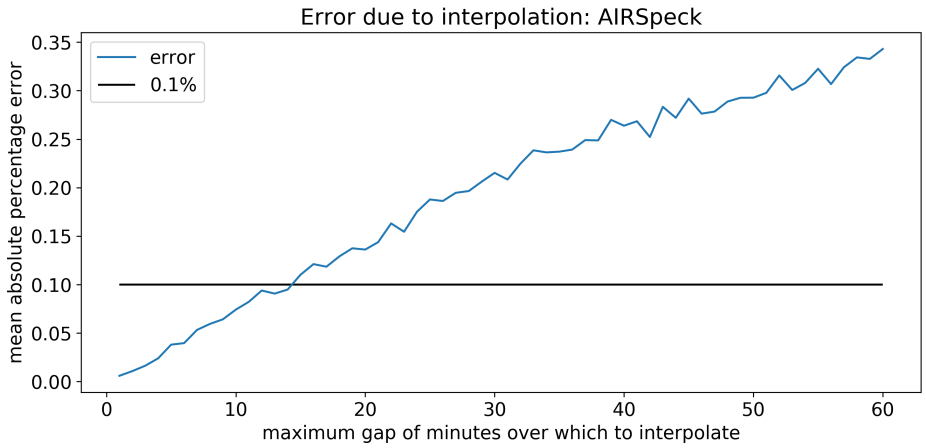

Figure 6.5: Percentage error of interpolation of AIRSpeck values

periment was repeated using linear, cubic spline and quadratic interpolation methods with linear interpolation performing best at all thresholds.

6.4 Standardising

Chapter 5 finds a strong association between respiratory rate and activity level during a given trial. Specifically, both the trend and volatility of a subject's respiratory rate reduce during periods of low activity e.g. at night while sleeping indoors. A correlation between personal exposure to $PM_{2.5}$ and time of day is also observed in Chapter 4: an increase in exposure in the morning is attributed to factors such as cooking in the house and traffic-related pollution. During the night, outdoor $PM_{2.5}$ levels increase due to an inversion layer. Such hourly patterns in both sets of data must be removed before performing any further statistical analysis as they can lead to biased results. This can be done by estimating local trend and volatility and standardising the data. The residual data would encode no information about activity level or time of day, leaving only other

drivers of respiratory rate and $PM_{2.5}$, with a possible causal relationship between the two.

The process of standardising a series of respiratory rate observations for an example trial (DAP022(2)) is now outlined. Standardising $PM_{2.5}$ data or any other series of data collected for a given trial is analogous[1].

Figure 6.6 plots observed respiratory rate for the trial. Variations in both trend and volatility over the change of day to night are clearly visible: around the hours of 23:30 to 05:30 on the first day of the trial and again around similar times on the second day there is a lower trend in respiratory rate as well as a lower variance in values (this can be identified by the plotted points appearing more bunched).

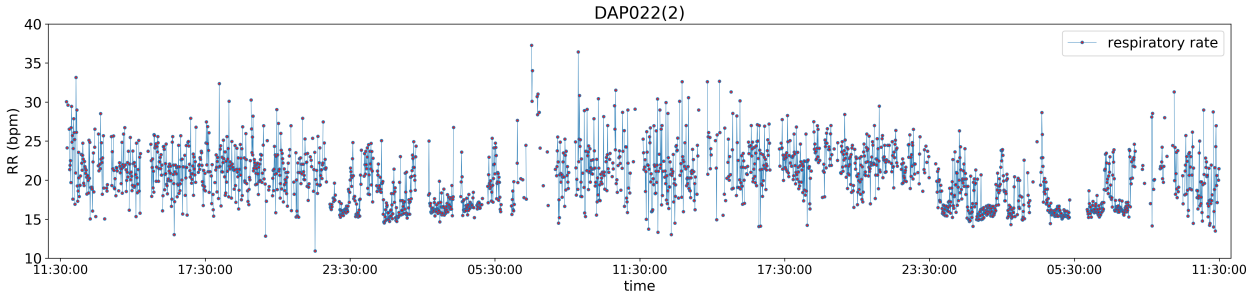

Figure 6.6: Respiratory rate observations for an example trial.

6.4.1 Local Trend Estimation

One of the simplest ways of estimating varying trend is through a moving average: for each observation x, a mean $\hat{x}$ is calculated as the average of the k nearest (in time) observations for some k. For the DAPHNE data, a moving average which weights all observations within a given window equally seems an oversimplification; a subject's respiratory rate should be closer in value to their respiratory rate in the next minute as opposed to the next five minutes. A moving average can be thought of as a simple type of *local regression* - one with a rectangular kernel. Some of the more standard types of local regression offer better ways of smoothing data and estimating trend.

Locally estimated scatterplot smoothing (LOESS) and locally weighted scatterplot smoothing (LOWESS) are two common and closely related *local* or *moving* regression techniques. This describes a family of non-parametric methods which combine standard regression with a nearest-neighbour approach in order to obtained smoothed series of data. LOESS is first discussed in [40] and has since been applied widely in the context of scatterplots. However it offers an elegant and relatively simple method for local trend estimation in time series. It is particularly applicable to settings with relatively sparse or noisy data, or settings in which traditional regression methods will fail.

[1] Analysis is carried out at the trial level therefore calculations for standardisation are made at the trial level.

Local regression fits a series of polynomials over windowed slices of a time series creating a smoothed fit of the data. The fitting at any given point x is weighted towards the data nearer to x. The size of window and degree of polynomial chosen introduce a standard bias-variance trade-off and hence must be selected carefully depending on the context. In addition it should be noted that the use of this technique assumes the local mean of a given observation can be estimated using polynomial regression, and that the estimation errors are white noise.

Python's *statsmodels* package offers an implementation of LOESS which can be fit to time series data. Figure 6.7 plots an estimated trend using a window size of 30 minutes and standard linear fits. The weighting function used in this implementation is a tricube function applied to the absolute distance between points - this is the same weighting function used in [40]. The differing trends from day to night are even more apparent in this plot and they can now be adjusted for by subtracting them from the original values. The adjusted respiratory rate values, now with zero mean, are shown in Figure 6.8.

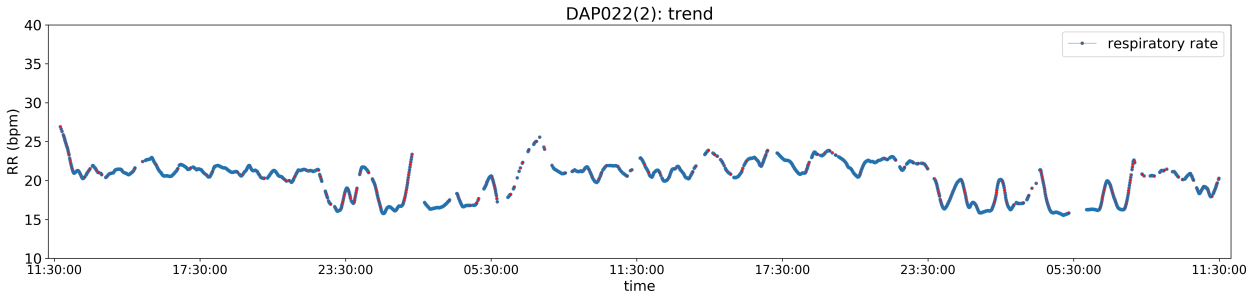

Figure 6.7: Respiratory rate estimated trend for trial DAP022(2) using LOESS

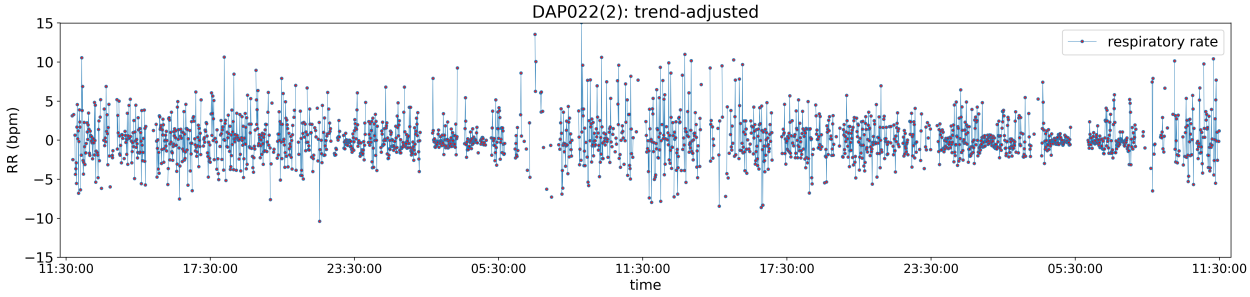

Figure 6.8: Trend-adjusted respiratory rate values for trial DAP022(2)

6.4.2 Volatility Estimation

The data in Figure 6.8 are stationary in mean but not in variance - the bunching of observations is still apparent during the hours of the night. Volatility is estimated by taking the standard deviation over a moving window - the chosen window size is 30 minutes in order to match the calculation of trend. The estimation of local volatility for trial DAP022(2) is shown in Figure 6.9. The final residual values are obtained from the adjusted series after dividing by volatility, these are shown in Figure 6.10.

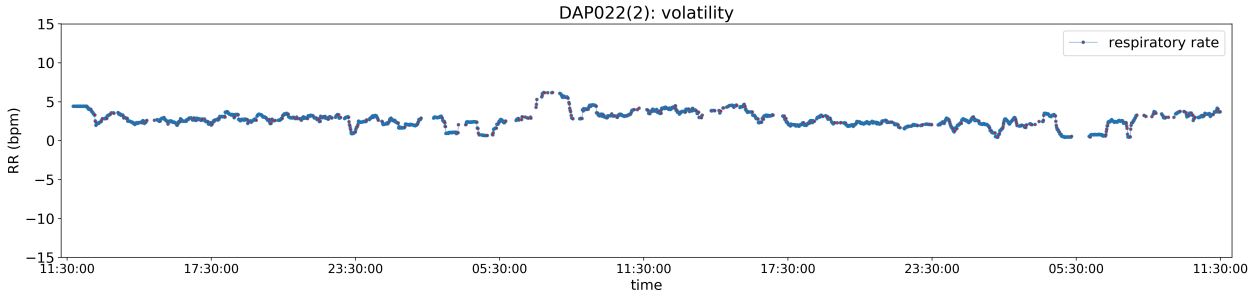

Figure 6.9: Estimated volatility of respiratory rate for trial DAP022(2)

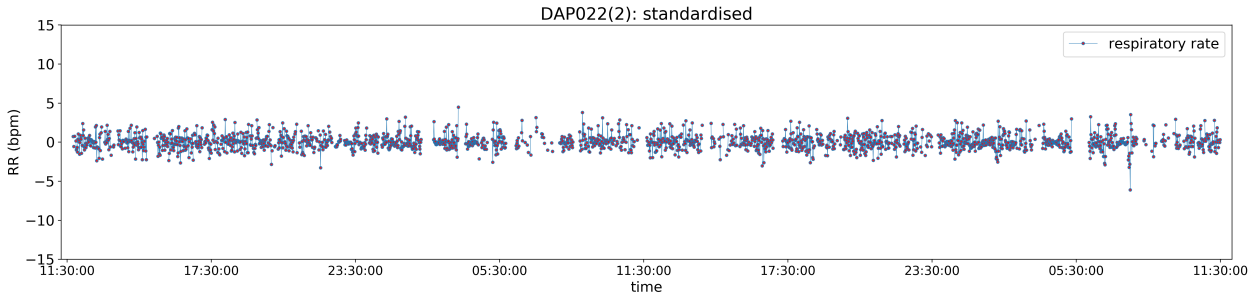

Figure 6.10: Standardised respiratory rate for trial DAP022(2)

To summarise, for each observation x local trend $\hat{x}$ and volatility $\hat{\sigma}_x$ are estimated. A residual $\tilde{x}$ is obtained by calculating $\tilde{x} = \frac{x-\hat{x}}{\hat{\sigma}_x}$.

Recall Figure 5.3 which plots the widely varying distributions in respiratory rate for a random subset of 50 trials. Figure 6.11 repeats the plot for a different random subset who have all been standardised as explained. Notice now that the distributions are all roughly the same - this makes future results comparable between subjects.

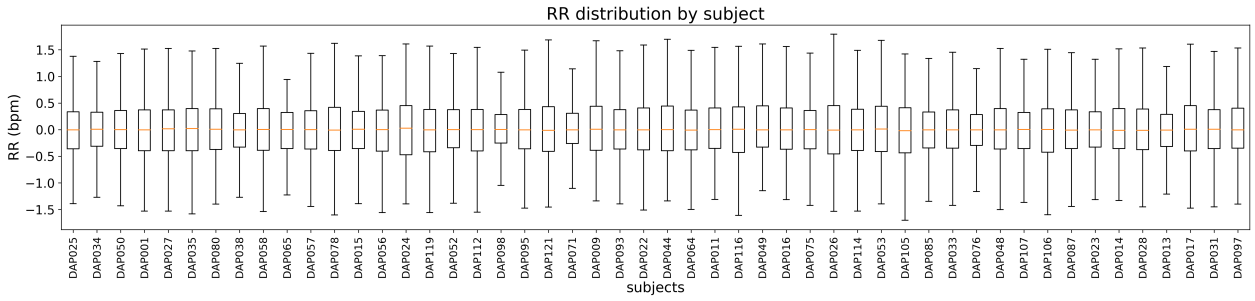

Figure 6.11: Distribution of respiratory rate observations for random subset of subjects (N=50) after local standardisation of trials.

6.5 Remaining Missing Data

In Section 6.3 gaps in RESpeck data of length up to 8 minutes are interpolated. For AIRSpeck data gaps of length up to 14 minutes are interpolated. Larger gaps in the data

do still exist as observed in Figure 6.3. Many implementations of statistical methods and tests do not offer masking of missing values and instead tend to drop them when executing. This is not appropriate for estimating exposure-response relationships as these depend on the absolute difference in time between observations - the lag. Dropping gaps in the data could lead to spurious results as an algorithm would interpret the observations on either side of the gap as occurring consecutively.

Remaining missing values are instead filled with the mean of the series. The standardised data has constant mean 0, hence missing values are filled with 0 value observations. While this solution is not ideal it will result in less potential for incorrect causal links to be identified than simply dropping the gaps altogether. Ideally, future trials should be carried out with extra care taken to obtain data without large gaps. Alternatively, more sophisticated methods of imputing the data could be explored, such as the use of recurrent neural networks for time series prediction, in order to increase the maximum gap over which missing values can be predicted.

Chapter 7

Initial Experiments

7.1 Partial Dependence Plots

The United States Environmental Protection Agency uses a qualitative and categorical system to evaluate causal relationships determined from observational data. This is criticised in [7] which instead favours treating such relationships as quantitative and continuous. In particular the use of partial dependence plots (PDP's) calculated from conditional probability tables is suggested as a much better alternative. These are used within the context of trained regression models and show how each predictor variable individually affects the model's predictions. In this sense they can be interpreted in the same way that the coefficients of a linear or logistic regression model are interpreted - showing how a change of one unit in a predictor variable results in a change of c in the prediction. A PDP plots this change in the prediction as a function of the change in the predictor. The calculation also marginalises over other predictor variables in the model - in this sense it isolates the effect of changing one variable. Existence of exposure-response relationships can be inferred from these plots by looking for an upward (or downward) trend. This section uses PDP's to investigate the effect of a change in personal $PM_{2.5}$ exposure on observed respiratory rate.

The data from all 183 trials are first pre-processed as described in Chapter 6, however, instead of filling large gaps with the mean these are dropped. This is because in the supervised learning problem each input-output pair is treated as a separate observation from the population being studied, hence the distance (in time) between different observations is not relevant and can be ignored. Any suitable regression model can be used - a gradient boosting machine is selected for this experiment to regress respiratory rate on $PM_{2.5}$, temperature and relative humidity. The latter two features are added to the model so they can be conditioned on when plotting the PDP as they are potential confounders. The model is trained on the entire dataset of 183 trials - while each subject is likely to respond differently to $PM_{2.5}$ exposure it is useful to first examine the overall effect of variations - this also increases the sample size for the model.

It may be that a change in $PM_{2.5}$ does not cause a change in respiratory rate immediately: there is a delay in the response. In order to observe how the underlying exposure-response relationship evolves over time a series of models can be trained which regress

respiratory rate on different lags (time shifted series) of the predictor variables. We train models and plot PDP's to investigate the immediate effect of changes in $PM_{2.5}$ as well as the delayed effect after 15, 30 and 60 minutes.

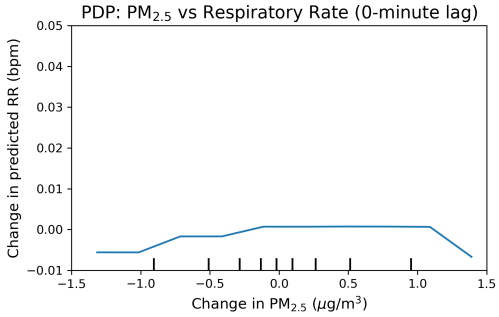

(a) Immediate response.

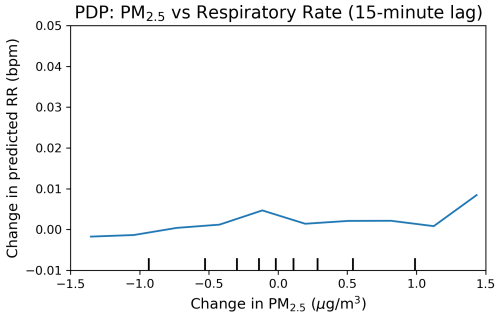

(b) Response after 15 minutes.

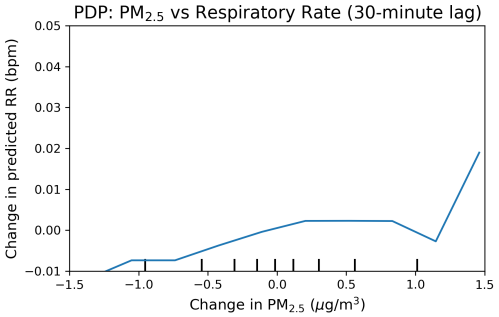

(c) Response after 30 minutes.

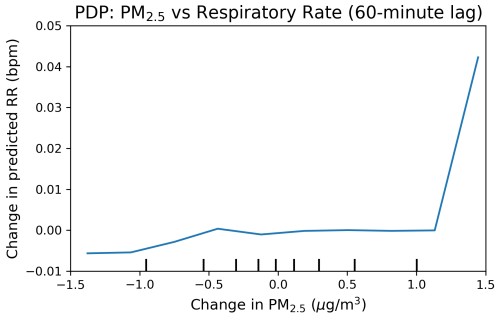

(d) Response after 60 minutes.

Figure 7.1: Partial dependence plots illustrating the response in predicted respiratory rate as a function of the change in $PM_{2.5}$ after different delays.

Figure 7.1 shows the 4 partial dependence plots illustrating the change in predicted respiratory rate as a function of change in $PM_{2.5}$ immediately and after periods of 15, 30 and 60 minutes (adjusted for temperature and humidity as well as activity level due to standardisation). The plots indicate that increases in $PM_{2.5}$ cause increases in respiratory rate and that the magnitude of the response is greater after a longer delay as the gradient of each plot in the interval 1.0-1.5 increases as the lag length increases. This suggests that on average, a large increase in $PM_{2.5}$ will cause an increase in a subject's respiratory rate and that this response will be most noticeable after a period of at least 1 hour.

While these results are clear evidence of exposure-response relationships between $PM_{2.5}$ and respiratory rate, further analysis is required to model these relationships accurately. Partial dependence plots are always produced in the context of a particular model which can only provide an approximation to the system being studied. In addition, the plots in Figure 7.1 are produced from a model trained on the entire dataset of 183 trials hence they only indicate a typical or average response to changes in $PM_{2.5}$. It is likely that at the subject level (and possibly even the individual trial level) the underlying exposure-response relationships will differ greatly.

7.2 Regression

The idea that $PM_{2.5}$ observations encode some information useful for predicting respiratory rate is not a notion of causality unless this information is unique. It is however a stronger idea than correlation alone which is symmetric: if X is correlated with Y then the converse is also true whereas the establishment that X is useful for predicting Y is a directional relationship (closer to causality). In this sense regression analysis can support the hypothesis of an exposure-response relationship existing.

Regression analysis can be carried out on the time series data from the DAPHNE study as follows: a model is constructed which regresses current respiratory rate observations on lagged observations. A second model is then constructed which adds lagged observations of personal $PM_{2.5}$ exposure to the input. If the second model improves on the performance of the first it is likely that $PM_{2.5}$ encodes useful information for predicting respiratory rate.

One of the simplest models used when describing stochastic processes is the autoregressive model. Given a time series X this assumes the current observation of X depends linearly on its own previous values and on one error term. An order p autoregressive process AR(p) is defined

$$X_t = c + \sum_{i=1}^{p} \varphi_i X_{t-i} + \varepsilon_t$$

where c is a constant, $\varphi_1, \ldots, \varphi_p$ are the parameters of the model and ε_t is a white noise term.

Notice that the parameters of the AR(p) model can be fit using standard linear regression making it an appropriate choice. This dissertation investigates the short-term effect of changes in $PM_{2.5}$ on respiratory rate and therefore we choose a relatively short lag period: $p = 60$ minutes.

An autoregressive process is a univariate case of the more general vector autoregressive (VAR) model which describes a multivariate time series dataset as a linear function of its previous observations. A VAR(p) process with k variables is defined

$$\begin{bmatrix} x_{1,t} \\ x_{2,t} \\ \vdots \\ x_{k,t} \end{bmatrix} = \begin{bmatrix} c_1 \\ c_2 \\ \vdots \\ c_k \end{bmatrix} + \begin{bmatrix} \varphi^1_{1,1} & \varphi^1_{1,2} & \cdots & \varphi^1_{1,k} \\ \varphi^1_{2,1} & \varphi^1_{2,2} & \cdots & \varphi^1_{2,k} \\ \vdots & \vdots & \ddots & \vdots \\ \varphi^1_{k,1} & \varphi^1_{k,2} & \cdots & \varphi^1_{k,k} \end{bmatrix} \begin{bmatrix} x_{1,t-1} \\ x_{2,t-1} \\ \vdots \\ x_{k,t-1} \end{bmatrix} + \cdots + \begin{bmatrix} \varphi^p_{1,1} & \varphi^p_{1,2} & \cdots & \varphi^p_{1,k} \\ \varphi^p_{2,1} & \varphi^p_{2,2} & \cdots & \varphi^p_{2,k} \\ \vdots & \vdots & \ddots & \vdots \\ \varphi^p_{k,1} & \varphi^p_{k,2} & \cdots & \varphi^p_{k,k} \end{bmatrix} \begin{bmatrix} x_{1,t-p} \\ x_{2,t-p} \\ \vdots \\ x_{k,t-p} \end{bmatrix} + \begin{bmatrix} \varepsilon_{1,t} \\ \varepsilon_{2,t} \\ \vdots \\ \varepsilon_{k,t} \end{bmatrix}$$

where $x_{i,t}$ is the observation of variable i at time t, c_i are constants, $\varphi^p_{i,k}$ are parameters of the model and $\varepsilon_{i,t}$ are white noise.

A VAR(60) process is used to describe observations of respiratory rate as a linear function of lagged observations of both itself and $PM_{2.5}$.

$$\begin{bmatrix} x_t \\ 0 \end{bmatrix} = \begin{bmatrix} c_1 \\ 0 \end{bmatrix} + \begin{bmatrix} \varphi_1^1 & \varphi_2^1 \\ 0 & 0 \end{bmatrix} \begin{bmatrix} x_{t-1} \\ y_{t-1} \end{bmatrix} + \ldots + \begin{bmatrix} \varphi_1^p & \varphi_2^p \\ 0 & 0 \end{bmatrix} \begin{bmatrix} x_{t-p} \\ y_{t-p} \end{bmatrix} + \begin{bmatrix} \varepsilon_t \\ 0 \end{bmatrix}$$

where x_t are respiratory rate values and y_t are $PM_{2.5}$ values. Again the parameters of this model can easily be fit using linear regression.

In practice any regression model could be used to evaluate the predictive power of $PM_{2.5}$. We use the following 3 models in our analysis:

- **Lasso Regression**: Linear regression models are prone to overfitting and poor generalisation. This manifests in needlessly large coefficients. L1 regularisation is a technique to avoid this by adding the absolute coefficient values to the error metric of the model (weighted by a free parameter) in order to restrict their magnitude and thereby reduce overfitting. Use of L1 regularisation is called lasso regression.
- **Elastic Net Regression**: L2 regularisation is similar to L1 in that it adds the squared coefficients to the error metric of the regression model. An elastic net combines both L1 and L2 regularisation.
- **Decision Tree Regression**: A decision tree used for regression seeks to break the training dataset down into smaller and smaller subsets in order to reduce the overall standard deviation of each subset. It is able to model a non-linear function unlike lasso regression and elastic net regression.

All models are implemented using Python's *scikit-learn* library ([53]). Using multiple models for this analysis helps to overcome individual limitations of any one particular model. For example both lasso regressors and elastic nets are unable to model non-linear associations unlike decision trees. The analysis is performed for each trial i.e. the two models, univariate and bivariate, are trained and evaluated for each trial. The results of the analysis are shown in Figure 7.2.

The results are identical for both lasso regression and elastic net regression while there are only a few instances of agreement with decision trees. There are large variations between trials highlighting the idea that no one exposure-response relationship will be able to describe every subject perfectly. It is now even more clear that investigating both linear and non-linear dependencies is a worthwhile endeavour as it may be that one exists but not the other. There are several trials for which none of the models improved in performance after the added input of $PM_{2.5}$. These should not be discarded as it is still possible that a very weak association between exposure and respiratory rate exists: one which is lost in the abstractions of the models used.

Figure 7.2: The results of the regression analysis explained in Section 7.2. The three models used are lasso regression (lasso), elastic net regression (en) and decision tree regression (cart). Green indicates an improvement in predictive performance using the bivariate model as opposed to the univariate model while red indicates a reduction in performance. The performance metric used is mean squared error.

Chapter 8

Granger Causality

Consider the following idea: suppose we assess our ability to predict a stochastic process Y given all the information in the universe. If we then construct a modified universe which excludes a stochastic process X and find our predictive power reduced, then we can conclude that X encodes some unique information about Y. This is the definition of Granger Causality as it is outlined in [29]. In the above example it is said that X *Granger causes* Y. Granger himself notes the impractical nature of having access to all the information in the universe, which is instead replaced by a set of observed values of both processes. It is then said that X Granger causes Y with respect to X although the latter phrase is usually omitted in literature, as is done hereafter.

Granger Causality has already been employed as a method of determining causal relationships in similar studies relating to air pollution. In [31] a standard parametric test for Granger Causality is used to establish a bidirectional causal relationship between the air quality in the cities of Beijing and Tianjin i.e. that pollution in neighbouring cities provably spreads. In [12] a GARCH model is used to describe influenza data in Taiwan, with a modified Granger Causality test based on this model finding that adult and elderly groups are strongly affected by increased $PM_{2.5}$ exposure.

8.1 Linear Granger Causality

8.1.1 The Experiment

Classical statistical methods of causal inference in time series data use parametric tests for Granger Causality, many of which have been implemented in data-processing languages like Python and R.

In particular, VAR-based tests assume a linear model for the data and look for linear relationships between the variables. A time series of respiratory rate values X is modelled as an AR(p) process:

$$X_t = c + \sum_{i=1}^{p} \varphi_i X_{t-i} + \varepsilon_t$$

This is referred to as the restricted model as it assumes respiratory rate only depends on past values of itself. The second model, referred to as the unrestricted model, is a VAR(p) process given by

$$X_t = c + \sum_{i=1}^{p} \alpha_i X_{t-i} + \sum_{i=1}^{p} \beta_i Y_{t-i} + \varepsilon_t$$

where in addition to respiratory rate X, personal $PM_{2.5}$ exposure Y is included. That is, the unrestricted model assumes dependence on past values of $PM_{2.5}$ as well as respiratory rate. The null hypothesis of the Granger Causality test is that the unrestricted model does not improve the predictive power of the restricted model (with rejection implying Y Granger causes X) $\rightarrow H_0$: $\forall i, \beta_i = 0$.

Under H_0 the Wald test statistic as calculated in [33] (7.6.3) asymptotically follows a χ^2 distribution, hence a standard χ^2 test of exclusion restrictions is suggested for Granger Causality. Unfortunately in practice the asymptotic distribution of the test statistic is often a poor approximation to the actual distribution due to small sample size, so an F-statistic is preferred (this can easily be calculated by dividing the Wald test statistic by its degrees of freedom - given by the lag length p). Python's *statsmodels* library provides an implementation of this exact test.

Tests for Granger Causality usually require stationary data (for the VAR-based test used here this is because the Wald test statistic only follows its usual asymptotic χ^2 distribution for stationary data). The DAPHNE data are pre-processed in order to be stationary in mean and variance using the methods described in Chapter 6. Stationarity is then verified through the Augmented Dickey-Fuller test: for a description of this test and further comments on stationarity refer to Section 3.2.2.

A linear relationship between observed respiratory rate values (measured by the RESpeck device) and observed exposure to $PM_{2.5}$ levels (measured by the personal AIRSpeck monitor) is tested for. A p-value is then calculated and checked for a rejection of H_0 at the 95% significance level. The test is carried out for each of the 183 trials and repeated for a series of increasing lag lengths p up to 1 hour - we choose 1, 5, 10, 15, 30, 45 and 60 minutes. The intuition behind this is that a rejected null hypothesis at a 30 minute lag versus a failure to reject at shorter lags suggests a delay of half an hour between exposure and response. Knowledge of such a delay can be useful in predicting potential asthma exacerbations. The results are shown in Figure 8.2 with a full table of p-values in Appendix A.

8.1.2 Results

Of the total 183 trials, 86 were found to show Granger Causality from $PM_{2.5}$ to respiratory rate at one or more of the lag lengths evaluated. This means 97 trials showed no evidence of Granger Causality. This is a large proportion of the data and there are many possible reasons for these results. Intuitively, it seems highly likely that both the restricted and unrestricted models used for this test are far too simple. When investigating complex systems in the real world it is usually inevitable that restricting assumptions will be made, however it is more reasonable to assume that relationships

between $PM_{2.5}$ and respiratory rate changes would be non-linear. It is also possible that some of the sample sizes of individual trials are too small - many are missing large chunks of data (as shown in Chapter 5) and while some of it has been imputed accurately, large gaps have to be filled with the mean (0) in order to carry out the test. This causes loss of information and can consequently lead to vague results.

Trials such as DAP120(1) and DAP073(2) find evidence of Granger Causality at all lag lengths tested. In fact there are 10 trials in total for which this is the case. Perhaps these subjects are particularly sensitive in their condition: showing signs of the effects of changes in $PM_{2.5}$ straight away with these effects lasting a long time. It would be beneficial to compare the general health of these 10 subjects and look for signs of poor respiratory condition as this would further support the results found here.

The test failed to report some or all p-values for a few trials such as DAP055(1) and DAP102(2). In the former case it was found that the ADF test failed to reject the null hypothesis of non-stationarity, meaning the test could not be carried out. In the latter, the implementation of the Granger Causality test failed to produce a result.

There is evidence of exposure-response relationships in trials such as DAP068(1), DAP041(1) an DAP028(1). Each shows a different kind of relationship:

- **DAP068(1)**: The p-values of the results of this trial are significant at lag lengths of 15, 30, 45 an 60 minutes. This suggests a delay between exposure and response during this trial of about 15 minutes.
- **DAP041(1)**: The p-values for this trial are significant at lag lengths of 5, 10, 15 and 30 minutes. This not only suggests that the delay between exposure and response during this trial for this subject is shorter than DAP068(1) (only around 5 minutes), it also suggests that the effects of exposure wear off over time and by 45 minutes after exposure the subject has returned to their regular pattern of breathing.
- **DAP028(1)**: The data for this trial shows one significant p-value when testing with a 1 minute lag length. This interesting result suggests that the subject feels the effects of exposure almost immediately but these effects also diminish just as quickly.

The above three examples illustrate the differences in response to $PM_{2.5}$ exposure between subjects. However we do see consistency between trials of the same subject. Figure 8.1 illustrates similar p-values observed in different trials from the same subject, indicating a consistency in their response to $PM_{2.5}$ exposure regardless of the time of year - this also indicates some consistency in health. A difference in results between trials of the same subject would not be too surprising however. Since different trials are sometimes months apart it is possible that the condition of a particular subject could either improve or deteriorate in that time.

While the VAR-based test for Granger Causality has produced valuable results it certainly still has its limitations. As already pointed out, the linear model assumed for the data is very restrictive and likely a poor fit. The experiment could be improved by using a test which is able to uncover non-linear relationships between the variables in

trial	1	5	10	15	30	45	60
DAP006(1)	0.50301	0.46739	0.38671	0.69306	0.00187	1E-05	7E-05
DAP006(2)	0.22263	0.2297	0.43924	0.03748	0.04645	0.02792	0.00401
DAP006(3)	0.43724	0.69488	0.0624	0.17192	0.02089	0.01963	7E-05
DAP055(3)	0.26559	0.04829	0.0243	6E-05	1E-05	0	3E-05
DAP055(2)	0.77813	0.00449	0.00211	1E-05	0.00016	8E-05	3E-05
DAP056(1)	0.67023	0.32626	0.32024	0.53819	0.38219	0.53525	0.28449
DAP056(2)	0.16406	0.35497	0.18322	0.17749	0.17752	0.2607	0.62351
DAP056(3)	0.58189	0.16379	0.39333	0.21774	0.2237	0.44598	0.30305
DAP070(1)	0.01689	9E-05	0.00041	0.00491	3E-05	5E-05	0.00085
DAP070(2)	0.00032	0.00489	0.00188	0.00101	0.00042	0.00209	1E-05
DAP074(1)	0.0579	0.00185	0.01733	0.04834	0.00462	0.00556	0.02702
DAP074(2)	0.1438	0.00818	0.00017	1E-05	0	0	0

Figure 8.1: Consistent exposure-response relationships for different trials by the same subject. Each column represents a Granger test at a different lag length (in minutes) and the value obtained is a p-value for Granger Causality. Significant p-values ($p < 0.05$) are coloured green, p-values between 0.05 and 0.1 are coloured yellow and p-values > 0.1 are coloured red.

the data. In addition, this test is unable to account for confounders: potential variables which drive changes in both $PM_{2.5}$ and respiratory rate. The results would be more reliable if a multivariate approach was taken in order to condition on confounders. The first of these issues is tackled in Section 8.2 and the second in Section 9.3.

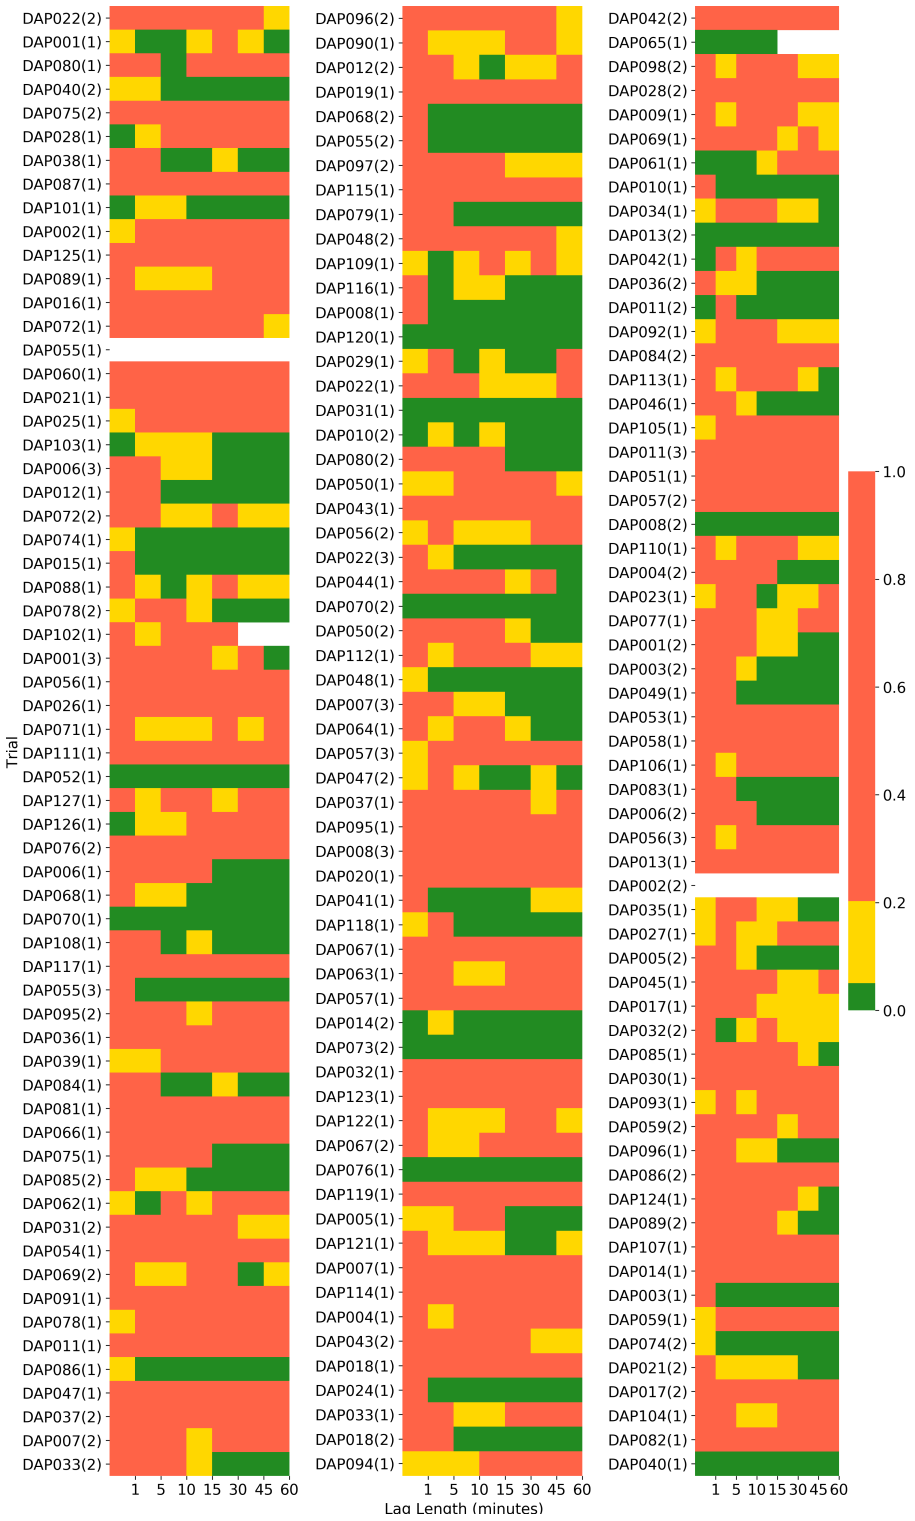

Figure 8.2: The results of the VAR-based tests for linear Granger Causality. Each trial is evaluated at lag lengths of 1, 5, 10, 15, 30, 45 and 60 minutes and a p-value for Granger Causality is obtained. Significant p-values ($p < 0.05$) are coloured green, p-values between 0.05 and 0.2 are coloured yellow and p-values > 0.2 are coloured red.

8.2 Non-Linear Granger Causality

8.2.1 The Experiment

In Section 8.1 the data are modelled by a VAR(p) process with linear relationships between the variables - this is a very strong assumption especially when modelling complex systems in the real world through observational data. The Granger Causality test used is unable to detect some non-linear relationships should they exist. Consider the following example model:

$$X_t = \alpha X_{t-1} Y_{t-1} + \varepsilon_t$$

where X and Y are i.i.d stochastic processes, ε is white noise and α is a parameter. The parametric test used in Section 8.1 would incorrectly fail to identify causality as all autocorrelations and cross-correlations are zero. Examples like this have motivated research into non-parametric tests for Granger Causality which can identify non-linear causal relationships.

General Granger Causality (sometimes referred to as non-linear Granger Causality) is expressed as a test for conditional independence. Formally, a time series X is said to Granger cause a time series Y if, for some $p \geq 1$

$$\mathbb{P}\left(X_{t+1} | X_t, \ldots, X_{t-p}, Y_t, \ldots, Y_{t-p}\right) \neq \mathbb{P}\left(X_{t+1} | X_t, \ldots, X_{t-p}\right)$$

Increased availability of greater computational power has made the task of testing this condition easier over the years with many methods being developed and implemented. The Hiemstra-Jones test, first applied in [30], is one of the most popular ones.

However in [8] the authors Diks and Panchenko illustrate how the test is unreliable. The relationship tested in the Hiemstra-Jones method is not actually implied by the null hypothesis of general Granger Causality, which can lead to the rejection rate tending to one at large sample sizes i.e. an increased chance of a type 1 statistical error. In [9], a modified test for non-linear Granger Causality is proposed and shown to be more stable.

The Diks-Panchenko test works as follows: we wish to investigate the existence of a causal relationship between two stochastic processes X and Y. As before we aim to do this by looking at the information held in lags of X and lags of Y since we cannot condition on the infinite past of the processes. Let $X_t^p = (X_t, \ldots, X_{t-p+1})$ and $Y_t^q = (Y_t, \ldots, Y_{t-q+1})$. The null hypothesis is that information from Y does not help in predicting X, which can be expressed as

$$H_0 : \mathbb{P}\left(X_{t+1} | X_t^p, Y_t^q\right) = \mathbb{P}\left(X_{t+1} | X_t^p\right)$$

which, on the condition that X and Y are stationary, is a statement about the invariant distribution of the $(p+q+1)$-dimensional vector (X_t^p, Y_t^q, Y_{t+1}). The null hypothesis describes conditional independence of the elements of this vector, therefore the joint probability density function $f_{X_t^p, Y_t^q, Y_{t+1}}\left(x_t^p, y_t^q, y_{t+1}\right)$ and its marginals must satisfy

$$\frac{f_{X_t^p, Y_t^q, Y_{t+1}}\left(x_t^p, y_t^q, y_{t+1}\right)}{f_{Y_t^q}\left(y_t^q\right)} = \frac{f_{X_t^p, Y_t^q}\left(x_t^p, y_t^q\right)}{f_{Y_t^q}\left(y_t^q\right)} \frac{f_{Y_t^q, Y_{t+1}}\left(y_t^q, y_{t+1}\right)}{f_{Y_t^q}\left(y_t^q\right)}$$

for each realisation of $\left(x_t^p, y_t^q, y_{t+1}\right)$. For the sake of clarity let us drop the indices and let $Z = Y_{t+1}$ so we can re-write the above as

$$\frac{f_{X,Y,Z}(x,y,z)}{f_Y(y)} = \frac{f_{X,Y}(x,y)}{f_Y(y)} \frac{f_{Y,Z}(y,z)}{f_Y(y)}$$

In [9] it is shown that H_0 implies

$$q_g \equiv \mathbb{E}\left[\left(\frac{f_{X,Y,Z}(X,Y,Z)}{f_Y(Y)} - \frac{f_{X,Y}(X,Y) f_{Y,Z}(Y,Z)}{f_Y(Y)}\right) g(X,Y,Z)\right] = 0$$

where $g(X,Y,Z)$ is a positive weight function. The authors choose $g(x,y,z) = f_Y^2(y)$ and derive an estimator for q_g:

$$T_n(\varepsilon) = \frac{(n-1)}{n(n-2)} \sum_i \left(\widehat{f}_{X,Y,Z}(X_i,Y_i,Z_i)\, \widehat{f}_Y(Y_i) - \widehat{f}_{X,Y}(X_i,Y_i)\, \widehat{f}_{Y,Z}(Y_i,Z_i)\right)$$

where

$$\widehat{f}_W(W_i) = \frac{(2\varepsilon)^{-d_W}}{n-1} \sum_{j, j \neq i} \mathbb{I}_{ij}^W$$

for $W = (X,Y,Z)$ and indicator function $\mathbb{I}_{ij}^W = \mathbb{I}\left(\left\|W_i - W_j\right\| < \varepsilon\right)$.

The bandwidth parameter ε is chosen such that it tends to zero at a rate proportional to the sample size; this causes the estimator to be consistent and avoids the problems of over-rejection which can occur with the Hiemstra-Jones test. For a sample size n the authors suggest $\varepsilon_n = \max\left(Cn^{-2/7}, 1.5\right)$ for a specific constant C which can be approximated by fitting two ARCH models to the series (in the bivariate case, for which [9] tackles) and averaging their respective coefficients (α_1).

Diks and Panchenko validate their test to prove its robust nature, making it a powerful tool for investigating non-linear Granger Causality. An implementation in C can be obtained from the authors on request and is used to investigate the existence of a non-linear causal relationship between changes in personal $PM_{2.5}$ exposure and changes in respiratory rate.

The main advantage of using a non-parametric test for causality is the ability to avoid making restricting assumptions on a particular model of our data. However, this test does have some issues which must be addressed before interpreting any results. At the time of writing, the C implementation can only be applied to a bivariate system meaning it is unable to account for common drivers of both $PM_{2.5}$ and respiratory rate. A rejection of the null hypothesis through this test can still be considered strong evidence for causality, but further analysis would be required before arriving at any definite conclusions. In addition, [9] notes in its evaluation of the test that the results are progressively conservative for increasing lag lengths. While not ideal, this is still preferable to rejections under the null hypothesis particularly within our domain.

The Diks-Panchenko test requires stationary time series. As mentioned in Section 3.2.2, while some non-parametric tests for stationarity do exist in literature they cover

a narrow range of situations and are not implemented. Hence non-linear stationarity is assumed here without proof.

The experiment in Section 8.1 is modified to apply the Diks-Panchenko test. Each of the 183 trials is tested for evidence of non-linear Granger Causality from $PM_{2.5}$ exposure to respiratory rate changes using the Diks-Panchenko test at 7 different lag lengths of 1, 5, 10, 15, 30, 45 and 60 minutes in order to uncover an exposure-response relationship. A p-value for each individual test is calculated and checked for a rejection of the null hypothesis of no causal relationship at the 95% significance level. The results are shown in Figure 8.3 with a full table of p-values in Appendix B.

8.2.2 Results

The Diks-Panchenko test has produced similar results to the VAR-based test used in Section 8.1. There are variations between subjects which indicate different sensitivity to $PM_{2.5}$ exposure and hence different exposure-response relationships. The differences between the two sets of results are now analysed.

The non-parametric test has highlighted new previously unidentified causal relationships. Specifically, 130 of the 183 trials have evidence to support the hypothesis of $PM_{2.5}$ Granger causing respiratory rate while the VAR-based test found such evidence in only 86 trials. This does still mean however that in 53 trials no evidence of a causal relationship was found. Note that it may well be that such relationships do exist but they are too weak for the test to identify them.

During 50 trials, tests for at least 6 of the 7 lag lengths chosen rejected the null hypothesis and concluded that $PM_{2.5}$ Granger causes respiratory rate, with 26 of these trials doing so at all 7 lag lengths of 1, 5, 10, 15, 30, 45 and 60 minutes. This is a large increase when compared with the figures from Section 8.1 and it shows the increased power of the non-linear test. In fact a total of 358 new causal relationships were found across all trials and all lag lengths. Trials such as DAP119(1) are particularly interesting as we now see evidence of Granger Causality at all lag lengths when previously none had been found whatsoever. This is an indication that while no linear relationship between $PM_{2.5}$ and respiratory rate exists for the data in this trial, a strong non-linear relationship does exist and has been identified.

The results for trials such as DAP005(1), DAP024(1) and DAP001(2) among others highlight something peculiar when compared to their counterpart figures in Section 8.1: the Diks-Panchenko test has failed to identify some causal relationships which were identified by the linear test. It is difficult to pinpoint the exact reason for this as there are multiple possibilities. One is that the relationship itself is linear and weak i.e. a narrow domain in which the previous approach may be more suited. This seems unlikely as there is no evidence to suggest that the Diks-Panchenko test lacks power in identifying linear relationships as a consequence of its increased generality. A more plausible reason for these changes is the conservative nature of the test under some conditions, first observed in [9]. One of its most attractive properties is that it controls for over-rejection of the null hypothesis; as a consequence it may sometimes under-reject at large sample sizes and large lag lengths. On inspection the vast majority of

instances where the test has failed to identify a linear relationship already discovered all occur at lag lengths beyond 30 minutes.

This experiment has helped to expand on the results of Section 8.1 to uncover a series of non-linear causal relationships between $PM_{2.5}$ and respiratory rate changes for the set of trials. However a key issue already brought up has not been dealt with: the potential for confounders to affect the test. Unfortunately the Diks-Panchenko test has only been implemented in the bivariate case - in Chapter 9 we make use of cutting-edge causal discovery methods which condition on potential confounders to uncover direct causal relationships.

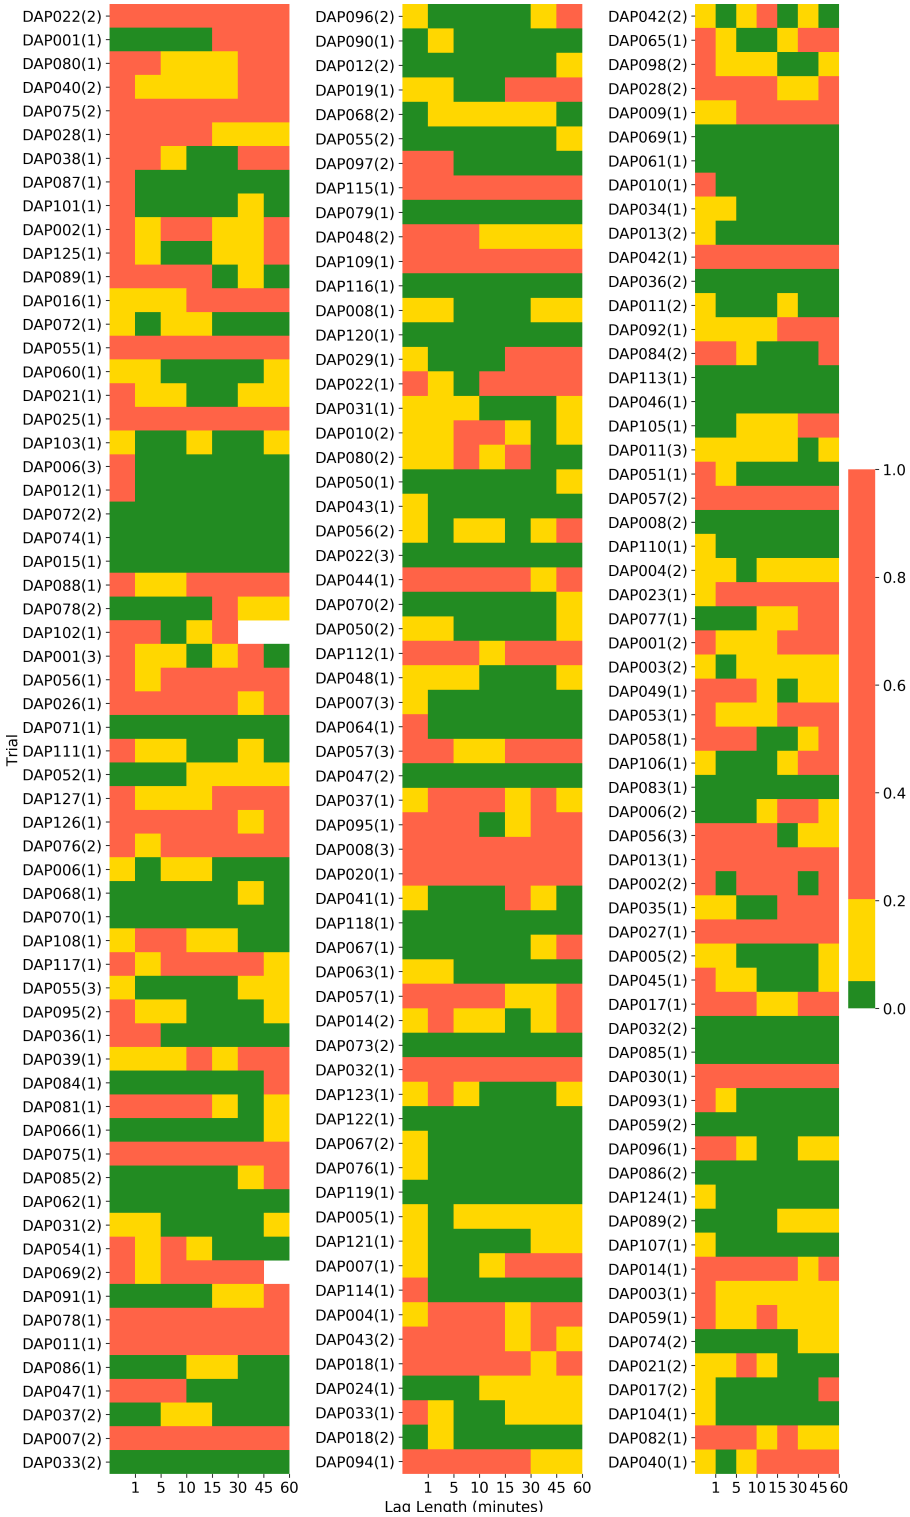

Figure 8.3: The results of the Diks-Panchenko tests for non-linear Granger Causality. Each trial is evaluated at lag lengths of 1, 5, 10, 15, 30, 45 and 60 minutes and a p-value for Granger Causality is obtained. Significant p-values ($p < 0.05$) are coloured green, p-values between 0.05 and 0.2 are coloured yellow and p-values > 0.2 are coloured red.

Chapter 9

Causal Discovery and Exposure-Response Relationships

The state of the art in causal inference leverages the advantages of access to big data and substantial computing power. PCMCI is a relatively new technique which achieves this goal. It is grounded in the paradigm of graphical causal models and causal discovery. In this chapter PCMCI is used to estimate exposure-response relationships from personal $PM_{2.5}$ exposure to respiratory rate.

9.1 Causal Networks

Graphical causal models construct a causal network: a directed acyclic graph (DAG) of all variables in a system at all time lags with edges indicating causal relationships between them. By nature they have no ambiguity between direct or indirect causal links - a useful tool when studying multivariate systems.

Consider a time series dataset of $N = 4$ variables $\mathbf{X} = \{X^1, X^2, X^3, X^4\}$. Figure 9.1 shows a possible causal network for the system underlying these data up to a maximum lag length of $\tau_{max} = 3$. There are multiple direct causal links which can be identified. Highlighted are:

- X^1_{t-2} and X^2_{t-2} directly cause X_{t-1} (red).
- X^2_{t-2}, X^3_{t-1} and X^4_{t-3} directly cause X^3_t (blue).

Given a variable X^i_t in a causal network, the direct predecessors of that variable are referred to as *causal parents*, with the set of all causal parents denoted $\mathcal{P}\left(X^i_t\right)$. Suppose an epidemiology study is investigating the effect of changes in a variable X^2 on another variable X^1. The subset of $\mathcal{P}\left(X^1_t\right)$ of lags of X^2 could be used to model an exposure-response relationship.

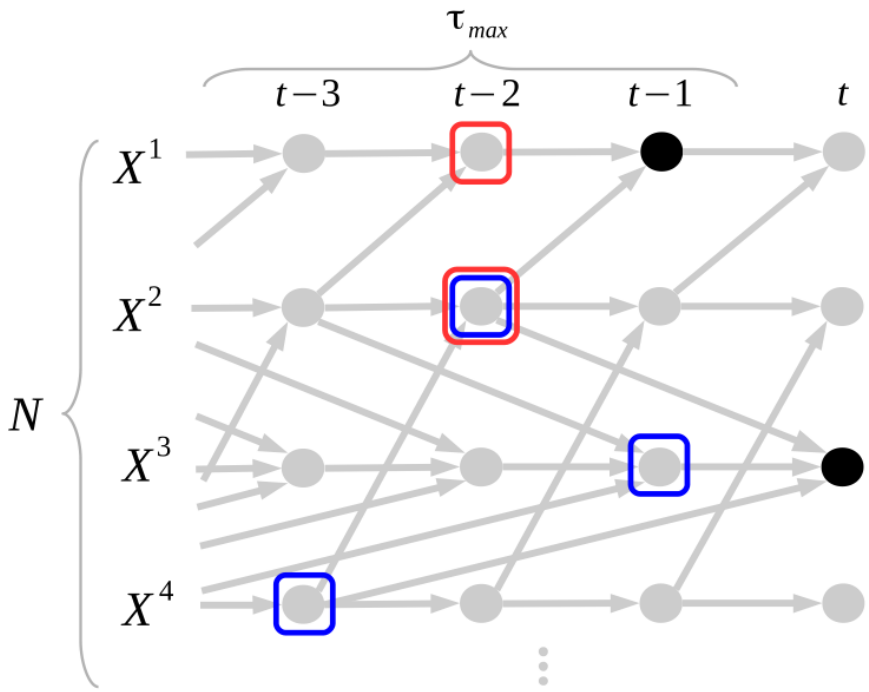

Figure 9.1: A section of an example causal network.

9.2 Causal Discovery

Specific tests for causality including Granger Causality tests are vulnerable to confirmation bias. Causal discovery overcomes this problem by following a different approach: the data are inspected ”at face value” and methods are employed to seek and identify any and all causal relationships in the underlying system. Many of these methods are also optimised for use with high-dimensional data - one of the issues with the experiments in Chapter 8 is that the tests are limited to bivariate systems (Granger Causality tests can quickly lose power at high dimensions - the curse of dimensionality). Examining multivariate systems with bivariate tests causes spurious results due to confounders.

The aim of causal discovery is to estimate all causal parents of each variable in a system from data: to estimate a causal network. This is non-trivial as the search space of possible DAGs is exponential in the number of variables, however it is usually reasonable to assume a sparse graph and this has led to the development of estimation methods which follow one of two paths:

- Initialise an empty causal network and recursively add edges.
- Initialise a fully connected undirected network and recursively remove edges.

The latter is the path taken in [44] where the authors Peter Spirtes and Clark Glymour propose the PC algorithm for DAG estimation which runs in polynomial time if the DAG to be estimated is sparse (exponential time in the worst case). Edges of the (initially) fully connected network are removed using a test for conditional independence - the reason for this is elaborated in Section 9.3. This test is left unspecified meaning

the algorithm can be used to search for both linear and non-linear dependencies by choosing one appropriately. A modified PC algorithm is used in PCMCI.

Before formally introducing PCMCI however, it is important to outline the necessary assumptions of causal discovery. [37] formally defines these assumptions along with examples to illustrate how they can be violated.

The first (and strongest) assumption of causal discovery is that of *Causal Sufficiency*. This requires that all common drivers (all variables which can cause two or more other variables in the system) are observed. In reality it is often difficult to account for every possible common driver without significant domain knowledge. In the context of the DAPHNE data, the AIRSpeck device records temperature and humidity, both of which are potential common drivers of $PM_{2.5}$ and respiratory rate. Therefore we include them in the system. This assumption also influences the maximum lag length evaluated and the sampling frequency of observations - this is discussed further in Section 10.2.

The second assumption is the *Causal Markov Condition* which requires that separation in a causal network implies independence (or the contrapositive that dependence implies an edge). This intuitively means that given the causal parents of a variable, all other information in the system is irrelevant for prediction. To meet this condition we assume that the internal conditions within the body which drive respiratory rate are encoded in recent past observations of it, while external drivers are measured ($PM_{2.5}$, temperature, humidity).

The *Faithfulness* assumption is the reverse implication of the Causal Markov Condition: conditional independence (a causal relationship) in the system implies an edge in the causal network. Together they imply that the causal network describes any and all causal relationships in the system.

Stationarity is a requirement for all time series data when applying causal discovery methods. Stationarity is discussed in Section 3.2.2 however it must be reiterated that non-parametric tests at the time of writing this dissertation are still narrow and lack implementations; they therefore cannot be carried out and stationarity is assumed without proof.

Causal discovery methods such as PCMCI can pair with any test for conditional independence. The test chosen may assume further qualities of the data e.g. only linear dependencies. The should be taken into account.

9.3 PCMCI

PCMCI is a causal discovery method which scales well with large multivariate time series and has strong false-positive control. Due to its use of the PC algorithm it can be used to estimate both linear and non-linear causal relationships. This is done by iterative testing for conditional independence. To see why, recall the practical definition of Granger Causality (Chapter 8): to say that X Granger causes Y means that past values of X contain unique information for predicting Y which is not contained in past values of Y alone. This is closely related to the more general concept of *transfer entropy* which measures the directed transfer of information between two processes i.e.

the reduced uncertainty in predicting Y given its past and the past of X. Formally,

$$T_{X\to Y} = H\left(Y_t|Y_{t-1:t-p}\right) - H\left(Y_t|Y_{t-1:t-p}, X_{t-1:t-p}\right)$$

Where $T_{X\to Y}$ denotes the transfer entropy from X to Y, Y_t denotes the current observation of Y, $Y_{t-1:t-p}$ and $X_{t-1:t-p}$ are the respective histories of Y and X up to lag p and H is Shannon's entropy (a measure of information). Transfer entropy is also written as *conditional mutual information* (CMI):

$$T_{X\to Y} = I\left(Y_t; X_{t-1:t-p}|Y_{t-1:t-p}\right)$$

Testing for causality is therefore testing whether the CMI above is nonzero. On inspection of the definition of CMI we see how this is equivalent to testing for conditional independence.

$$\begin{aligned} I(X;Y|Z) = \iiint \mathbb{P}(x,y,z)\log\frac{\mathbb{P}(x,y|z)}{\mathbb{P}(x|z)\cdot\mathbb{P}(y|z)}\,dx\,dy\,dx = 0 \\ \iff \mathbb{P}(x,y|z) = \mathbb{P}(x|z)\cdot\mathbb{P}(y|z) \quad \forall x, y, z \\ \iff X \perp Y|Z \end{aligned}$$

With this understanding in place we can now formally define a causal relationship in a causal network. A causal links exists from variables $X^i_{t-p} \to X^j_t$ if they are *not* conditionally independent given the past of the whole system $\mathbf{X}$. Note that instead of conditioning on the past of the entire process, conditioning on the set $\mathcal{P}(X^j_t)\backslash\{X^i_{t-p}\}$ would be sufficient due to the Causal Markov Condition.

9.3.1 The Method

PCMCI is a two-stage process. The first stage (the PC stage) is a Markov set discovery algorithm used to estimate a superset of the causal parents of every variable in the system through a modified version of the PC algorithm. The pseudocode is given in Algorithm 2 with a brief explanation below. For more detailed pseudocode and a complete description refer to [15].

The intuitions motivating Algorithm 2 are as follows:

- The algorithm begins by initialising a superset of all possible causal parents of X^i_t.
- The outer loop repeats for all possible subset sizes of $\hat{\mathcal{P}}\left(X^i_t\right)$. Note that $|\hat{\mathcal{P}}\left(X^i_t\right)| = N\cdot$ max_lag.
- PCMCI is fast because it sorts $\hat{\mathcal{P}}\left(X^i_t\right)$ based on the chosen statistic and only tests the q parents with strongest dependency instead of every possible subset $\mathcal{S}$ with $|\mathcal{S}| = q$. The test itself is for conditional independence - if the null hypothesis cannot be rejected at an appropriate significance level the link can be removed from $\hat{\mathcal{P}}\left(X^i_t\right)$.
- In this sense the algorithm recursively removes links it is confident enough to remove in order to obtain the superset $\hat{\mathcal{P}}\left(X^i_t\right)$ of the causal parents of X^i_t.

Algorithm 2: PCMCI: First Stage
The data are considered a time series dataset $\mathbf{X} = \left(X^1, X^2, \ldots, X^N\right)$ of N variables.

input : variable X_t^i to evaluate, maximum lag length max_lag, conditional independence test function
output: superset of causal parents of X_t^i, $\hat{\mathcal{P}}\left(X_t^i\right)$

$\hat{\mathcal{P}}\left(X_t^i\right) = \{\mathbf{X}_{t-1}, \mathbf{X}_{t-2}, \ldots, \mathbf{X}_{t-\text{max\_lag}}\}$;
for $q \leftarrow 0$ **to** $|\hat{\mathcal{P}}\left(X_t^i\right)|$ **do**
  sort $\hat{\mathcal{P}}\left(X_t^i\right)$ based on chosen test statistic;
  $\mathcal{S} =$ first q parents in $\hat{\mathcal{P}}\left(X_t^i\right)$ based on sorting;
  for $X_{t-k}^j \in \hat{\mathcal{P}}\left(X_t^i\right)$ **do**
    test $H_0 : X_{t-k}^j \not\perp X_t^i | \mathcal{S}$;
    if H_0 *is not rejected* **then**
      $\hat{\mathcal{P}}\left(X_t^i\right) = \hat{\mathcal{P}}\left(X_t^i\right) \backslash \{X_{t-k}^j\}$;
return $\hat{\mathcal{P}}\left(X_t^i\right)$;

Algorithm 2 is repeated for every variable $X_t^i \in \mathbf{X}_t$.

In the second stage each remaining link is tested for *momentary conditional independence* (MCI). This is an extension of the standard test for conditional independence which conditions on the causal parents of *both* variables in the link. That is, for a link $X_{t-p}^i \to X_t^j$ the MCI test is

$$X_{t-p}^i \perp X_t^j | \hat{\mathcal{P}}\left(X_t^j\right) \setminus \{X_{t-p}^i\}, \hat{\mathcal{P}}\left(X_{t-p}^i\right).$$

The additional conditioning controls for highly autocorrelated data and leads to much lower false-positive rates than other commonly used causal discovery methods. The significance of each link can finally be assessed from the p-values of this test.

PCMCI is used to uncover causal relationships in the multivariate system of respiratory rate, $PM_{2.5}$, temperature and humidity for a subset of 59 of the original 183 trials. This is done to reduce time taken running experiments as the method, while faster than other causal discovery methods, is computationally expensive. The 59 trials are chosen based on missing data - these are trials for which less than 40% of observations across all 4 observed time series are missing. The existence of both linear and non-linear dependencies is investigated using two different tests for conditional independence. The method (and tests) are implemented using the open-source *Tigramite* ([52]) Python package maintained by Runge himself. The data are pre-processed using the techniques described in Chapter 6.

For linear dependencies the test for conditional independency $X \perp Y | Z$ is carried out using partial correlation. This is done by fitting two multivariate regression models to predict both X and Y given Z. The correlation of the residuals is then evaluated, with test assuming these residuals are approximately normally distributed. All linear

dependencies up to a maximum lag length of 60 minutes are evaluated. In addition, Tigramite offers methods for ignoring missing data and masking certain observations. This is used to repeat the test for each trial at the maximum lag length possible while ignoring all missing data and masking to isolate instances of increases in $PM_{2.5}$ from one minute to the next.

The test used for non-linear dependencies is a fully non-parametric test based on conditional mutual information (recall that conditional independence implies this is zero). The test is fully described in [38] and briefly explained here. The following estimator for conditional mutual information is used:

$$\widehat{I}(X;Y|Z) = \psi(k) + \frac{1}{n}\sum_{i=1}^{n}\left[\psi\left(k_i^z\right) - \psi\left(k_i^{xz}\right) - \psi\left(k_i^{yz}\right)\right]$$

where ψ is the Digamma function $\psi(x) = \frac{d}{dx}\ln\Gamma(x)$, n is the sample length and k is a parameter specifying the number of nearest neighbours of each sample to be taken from the joint sample space. The main advantage of this nearest-neighbour estimator approach is that it can capture almost any non-linear dependency - in this sense the test is fully non-parametric. Note however that three separate nearest-neighbour calculations have to be carried out for each sample. This drastically increases runtime to an infeasible time-frame so the causal network is instead approximated using only lags at lengths 1, 5, 10, 15, 30, 45 and 60 minutes (matching the lag lengths used in the Granger Causality tests in Chapter 8).

9.3.2 Results

To reiterate, linear dependencies from $PM_{2.5}$ to respiratory rate for the subset of 59 trials are investigated using PCMCI with a partial correlation test - the results are shown in Figure 9.4. For a full table of p-values see Appendix C. A total of only 186 significant causal links are found across all trials, meaning each trial has an average of about three. PCMCI is a powerful method with strong false-positive control, especially when compared to other tests for linear dependencies such as the VAR-based test for Granger Causality. With this in mind 186 causal links is a significant result as it is evidence that the exposure-response relationships have linear components. In fact, there are only two trials for which the method finds no linear dependencies whatsoever: DAP014(2) and DAP056(3).

Figure 9.2 is a histogram of the number of linear causal relationships found at each lag length (across all 59 trials). It approximates the distribution of causal relationships with any peaks giving an indication of lag lengths for which a link is commonly found. There is a clear peak around 15-25 minutes which shows that the typical delay between exposure and response is quite short. There is a hint of a second smaller peak around the 45 minute mark however this is difficult to verify without repeating the experiment on a larger dataset. If this were to exist it might suggest that the body's typical response to exposure occurs in waves which may corresponding to the level of penetration of the particulate matter into the respiratory system.

These results can be compared with those of Section 8.1 for the VAR-based test for Granger Causality. It is clear that the results of PCMCI are much more conservative in

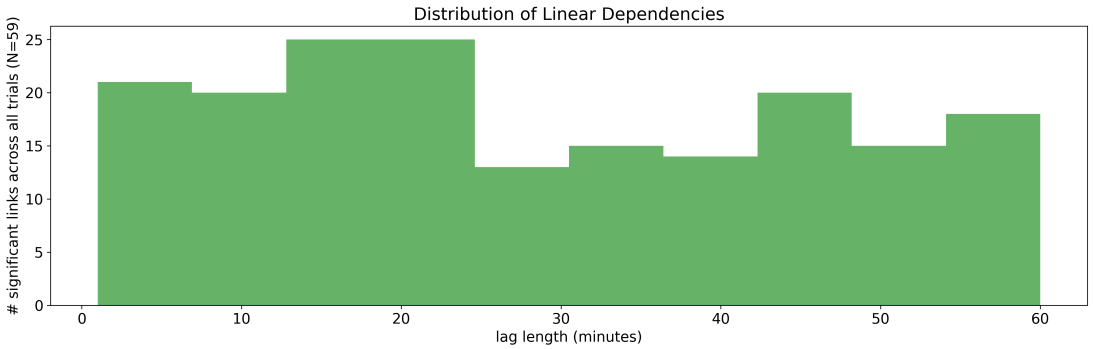

Figure 9.2: Number of occurrences of causal links at each lag length across all trials tested

their identification of causal links and this can be attributed to two main reasons. It is likely that some of the evidence of Granger Causality found in Section 8.1 is incorrect due to confounders such as temperature and humidity. The Granger test is applied to the bivariate case whereas PCMCI works with multivariate systems, conditioning direct links between $PM_{2.5}$ and respiratory rate on the other variables. It is also possible that false-positives due to strong autocorrelations in the variables occur in the results from Section 8.1. The MCI stage of PCMCI controls for this and therefore the new results are more trustworthy.

Many time series in the DAPHNE data have large gaps, sometimes hours long due to subjects removing their sensors during the night (Section 6.5). The significant proportion of missing data is problematic as most implementations of statistical tests and methods for causal inference are unable to deal with gaps. Due to this we opt to interpolate as much missing data as possible. After local standardisation to zero mean and unit variance, imputation of any remaining gaps simply uses the mean. While this is a partial solution it results in loss of information and in cases where most of the data are missing, potentially spurious results. The Tigramite library is able to consistently handle missing values. It does this by dismissing all time slices where missing values occur while still dealing with time lags correctly. Of course, this causes a reduction in sample size which scales with increasingly sparse data and in some cases can restrict the maximum lag length used in the test. We repeat the use of PCMCI assuming linear dependencies for each of the 59 trials, running the algorithm for the largest possible lag length which can be used without inaccuracy due to lack of samples.

Tigramite offers another powerful function: masking of values. This can be used to include or exclude samples depending on some conditions, for example when analysing yearly climate data it is possible to focus on the causal parents of a target variable only in winter months. In the context of the DAPHNE data, while we wish to investigate exposure-response relationships from $PM_{2.5}$ exposure to respiratory rate changes it would be useful to isolate the effect of *increases* in $PM_{2.5}$. This is done by marking all $PM_{2.5}$ observations which are an increase from the previous minute.

We carry out a modified experiment to examine linear dependencies from *increases* in $PM_{2.5}$ to respiratory rate, ignoring missing data. The results are shown in Figure

9.5 with a table of p-values also in Appendix C. The p-values differ from the previous experiment due to the examination of a different causal relationship (the effect of *increases* in $PM_{2.5}$ on respiratory rate). Adjusting for the number of samples available in each trial when testing at a particular lag length, a new distribution of causal links is plotted in Figure 9.3. Small sample size results in this distribution covering only half the range of Figure 9.2 however the results are still interesting. There are two clear peaks; one around 5 minutes which tapers off gently and another steeper peak around 25-30 minutes. This shows that increases in $PM_{2.5}$ usually cause a very quick response in respiratory rate (within 5 minutes) followed by another response after a longer period of time (around 30 minutes), possibly after the particulate matter has had time to enter deeper into the body[1].

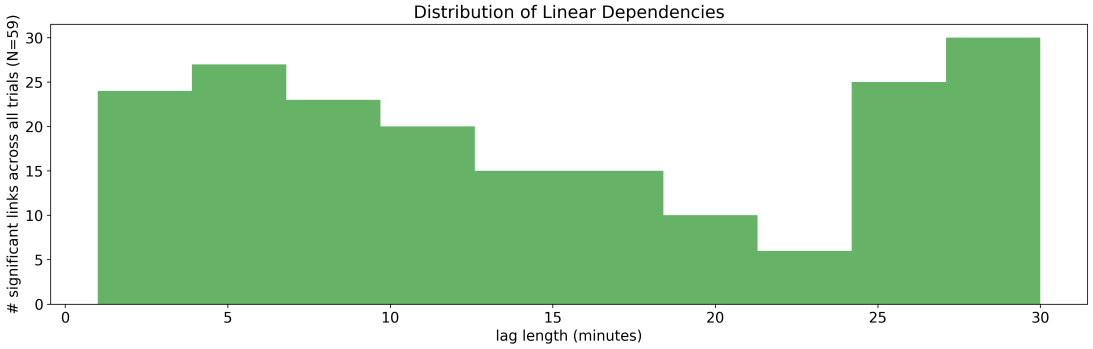

Figure 9.3: Distribution of linear causal links across all lag lengths and all trials

PCMCI is also used to investigate non-linear causal relationships in the data. This is done with a fully non-parametric test which uses a nearest-neighbour-based estimator for conditional mutual information and is therefore able to capture almost any type of dependency. The results are shown in Figure 9.6 with a full table of p-values in Appendix D. This test is considerably more computationally expensive than the linear test and is therefore carried out for lags of 1, 5, 10, 15, 30, 45 and 60 minutes only. The results from this test are overwhelmingly in favour of causality, in fact over half of the p-values for causality in Figure 9.6 are significant: 237 of them. The mean number of causal links for each trial evaluated at the 7 lag lengths chosen is 4 and only 7 trials showed no evidence of causality at all. This shows that in the majority of cases, a subject's respiratory rate is affected by changes in $PM_{2.5}$.

18 trials had direct causal links between $PM_{2.5}$ and respiratory rate for at least 6 of the 7 lag lengths tested. For these trials, response to changes in exposure is within 5 minutes if not immediate and persists for at least a full hour. This is motivation to repeat the experiment with a longer maximum lag length in order to determine the full length of the exposure-response relationship.

Table 9.1 shows the number of significant causal links found at each lag length tested across all 59 trials. It is clear that it is more common to find causal dependencies at earlier lags from around 5 minutes until somewhere just under half an hour. This roughly

[1]Note that these responses are characterised by *linear* dependencies. Non-linear relationships are not examined in this analysis.

Table 9.1: Number of trials with significant causal links at different lag lengths (test for non-linear dependencies).

Delay in Minutes	**1**	**5**	**10**	**15**	**30**	**45**	**60**
# Causal Links (N=59)	22	47	45	42	32	28	21

coincides with the results from the tests for linear dependencies with the general result being that the exposure-response relationship from $PM_{2.5}$ to respiratory rate typically peaks in that short timeframe.

For the most part, the results shown in Figure 9.6 differ very little from the results of the Diks-Panchenko test for non-linear Granger Causality (Section 8.2). Figure 9.7 shows any corresponding p-values which differ in their significance ($p < 0.05$) between the tests. In general, the results show that where differences are found, new causal links were more commonly discovered at earlier lags by PCMCI and discarded at later lags. In the first situation it is likely that PCMCI is able to uncover more complex relationships due to the flexibility of the nearest-neighbour-estimator for conditional mutual information. This allows it to identify relationships in the data which remain hidden from the Diks-Panchenko test. The latter case is particularly interesting as the Diks-Panchenko test is known to be conservative at larger lag lengths. It is likely then that the causal relationships misidentified by the Diks-Panchenko test are false-positives due to confounding or residual autocorrelation, both of which have been controlled for in PCMCI.

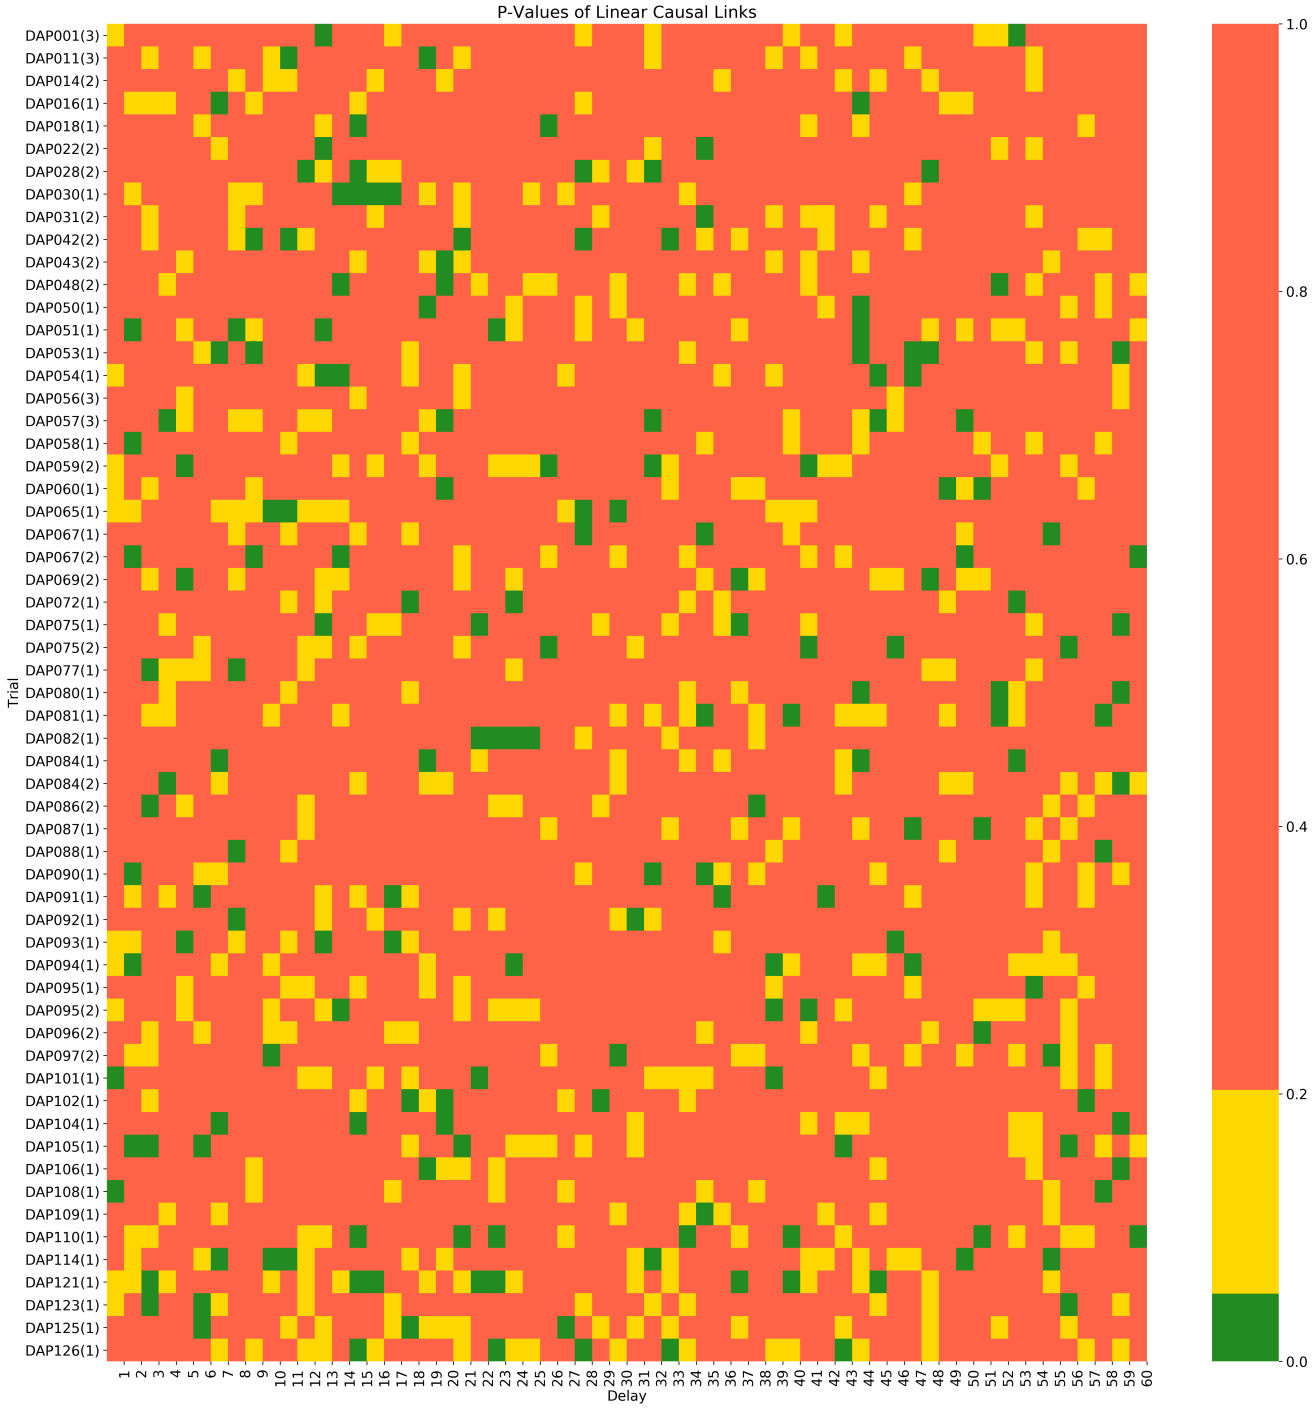

Figure 9.4: Results from PCMCI test for linear dependencies from $PM_{2.5}$ to respiratory rate for each trial for lag lengths from 1 to 60 minutes. Values are p-values of the MCI test. Significant p-values ($p < 0.05$) are coloured green, p-values between 0.05 and 0.2 are coloured yellow and p-values > 0.2 are coloured red.

Figure 9.5: Results from PCMCI test for linear dependencies from $PM_{2.5}$ increases to respiratory rate for each trial. Missing values are ignored which limits the maximum lag length which can be evaluated due to small sample size. Values are p-values of the MCI test. Significant p-values ($p < 0.05$) are coloured green, p-values between 0.05 and 0.2 are coloured yellow and p-values > 0.2 are coloured red.

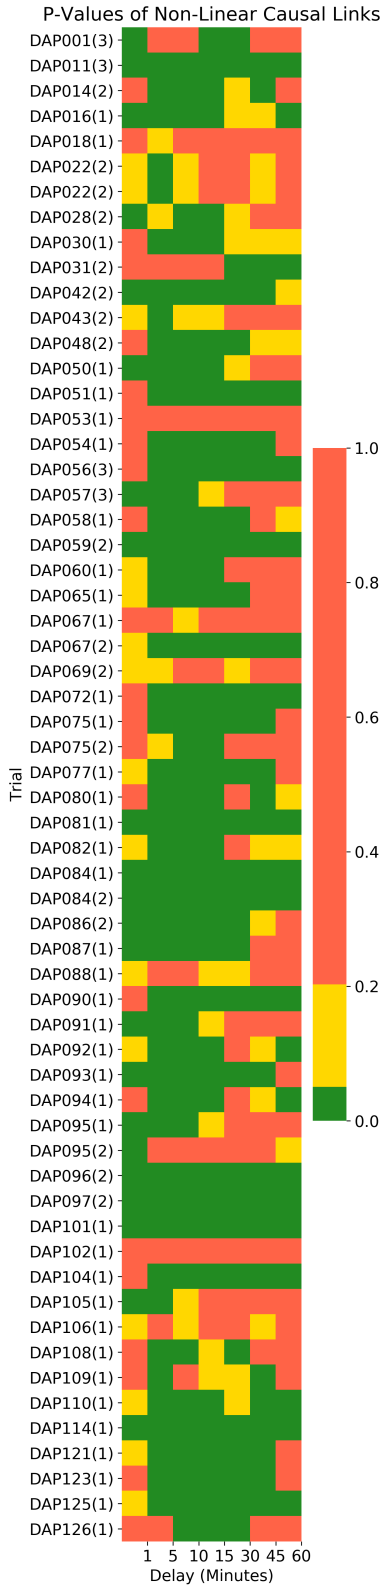

Figure 9.6: Results from PCMCI test for non-linear dependencies from $PM_{2.5}$ to respiratory rate for each trial for lag lengths 1, 5, 10, 15, 30, 45 and 60 minutes. Values are p-values of the MCI test. Significant p-values ($p < 0.05$) are coloured green, p-values between 0.05 and 0.2 are coloured yellow and p-values > 0.2 are coloured red.

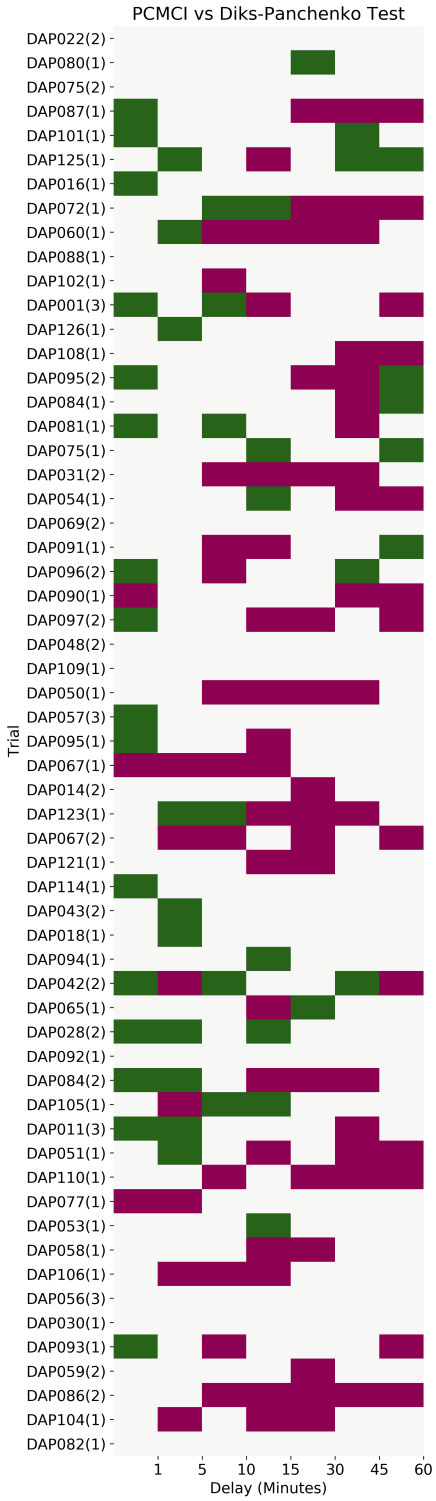

Figure 9.7: Changes in the results from PCMCI versus the Diks-Panchenko Test for non-linear Granger Causality. Green indicates a causal link identified by PCMCI which was unidentified by Diks-Panchenko. Pink indicates a causal link identified by Diks-Panchenko which was rejected by PCMCI.

Chapter 10

Conclusion

10.1 Discussion and Strengths

In this dissertation we analysed data from 183 trials of the DAPHNE study to understand trends and associations in air pollution exposure and respiratory rate for asthmatic adolescents in Delhi. We aimed to estimate exposure-response relationships from personal $PM_{2.5}$ exposure to changes in respiratory rate for each of these trials using two methods of causal inference: Granger Causality and PCMCI. We chose to focus on short-term ($<$ 1 hour) effects as these have not been examined in great depth before.

The data were first cleaned and pre-processed. After observing hourly trends in both respiratory rate (due to activity level) and $PM_{2.5}$ (due to nightly inversion layers in the atmosphere) we locally standardised all trials to zero mean and unit standard deviation in order to remove these trends and ensure that results between different subjects were comparable.

10.1.1 Granger Causality

One of the key strengths of this project was the exhaustive search of all types of causal relationships. We used both parametric and non-parametric tests for Granger Causality to investigate both linear and non-linear relationships in the data. In the former case we made use of the well-known VAR-based test for Granger Causality. Our approach of testing at multiple lag lengths (up to a maximum of one hour) in order to better gauge the delay between exposure and response was well thought out and new, although in this particular case the results were mixed: the test found evidence of Granger causality in only a minority of trials such as DAP073(2) where significant causal links were found at all lag lengths tested or DAP049(1) where Granger causality was observed after 5 minutes (these results are seen in Figure 8.2). The low number of significant results was attributed to the rather limiting linear model for the data. It seemed more intuitive for the human body to have a non-linear relationship with air pollution exposure. We tested for such a relationship using the more powerful Diks-Panchenko test - this was chosen after careful consideration of the most common non-parametric tests

for Granger Causality. We identified exposure-response relationships in several trials as shown in Figure 8.3, including many new relationships unidentified in the linear test; DAP119(1) is one example of this. For these trials, evidence of Granger causality was found to be significant at the 95% level - often at multiple lag lengths. The results not only suggested that response to exposure is rapid, they also hinted that the lasting effects can have a longer duration than one hour, motivating future work to increase the maximum lag length used in these analyses.

10.1.2 PCMCI

We sought to make use of computationally expensive causal discovery methods which learn from big data. We specifically chose PCMCI, a new method proven to be very powerful with sophisticated false-positive control and the ability to deal with multivariate time series data accurately. This allowed us to account for potential confounders such as temperature and relative humidity. We used PCMCI on a subset of 59 trials: trials for which less than 40% of observations across all 4 observed time series (respiratory rate, $PM_{2.5}$, temperature and relative humidity) were missing. PCMCI was first paired with a partial correlation test statistic to search for linear dependencies in the DAPHNE data. We found infrequent yet significant causal relationships between lags of $PM_{2.5}$ and respiratory rate, with these relationships occurring more commonly at earlier lags generally under half an hour (Figure 9.4).

Using a convenient feature of the implementation of PCMI in the Tigramite Python package, we were able to repeat the test and mask missing values in the data instead of imputing them with the mean - this was done for the same subset of 59 trials as the modified experiment was unstable for trials with too much missing data. In this experiment we also isolated the effects of increases in $PM_{2.5}$, finding that these effects can occur on an even shorter timescale: within 5 minutes of exposure as shown in Figure 9.5.

Finally, we made use of a fully non-parametric test for causality which utilised a nearest-neighbour based estimator of conditional mutual information. This flexible test was able to discover almost any type of non-linear relationship in the data. The test was run at a series of 7 lag lengths over the course of an hour as a full experiment over all 60 minutes of lags as before would take weeks of CPU time to run. However, as shown in Figure 9.6, we found evidence significant at the 95% level for non-linear direct causal relationships between $PM_{2.5}$ and respiratory rate for almost every trial tested, with the results indicating that response most commonly occurs in the period of 5-15 minutes after exposure. Thus we have shown that $PM_{2.5}$ causes short-term (< 1 hour) changes in respiratory rate, adjusting for the confounders of activity level, temperature and relative humidity. This is the first study of its kind to demonstrate this direct dependence from observational data. Since previous research such as [10], [27] and [46] find positive associations between increases in $PM_{2.5}$ exposure and increased respiratory rate, our results suggest that increases in $PM_{2.5}$ can *directly cause* increases in respiratory rate.

10.2 Limits, Improvements and Future Work

The significant gaps in the data caused problems in all experiments carried out in this project. Our partial solution of interpolating as much data as possible and filling any remaining gaps with the mean was not ideal as it resulted in loss of information. While Tigramite's functionality to mask missing values offered a complete solution, the resulting loss in sample size was sometimes too great to run the algorithm. This problem could be mitigated by improving the accuracy of imputation of missing values, perhaps through the use of recurrent neural networks which have shown much promise in this domain.

Our experiment using PCMCI to identify non-linear dependencies in the data has scope for improvement. Due to time-constraints, the estimated causal network had to be simplified by reducing the number of lags tested. This has the potential to miss relationships in the data by violation of the assumption of Causal Sufficiency. Future work could devote the significant time and resources required to test for all lag lengths up to an hour or even beyond to a full day. The method could also be extended to learn from multiple samples instead of one time series dataset, allowing it to estimate causal networks for each subject by learning from multiple trials.

This dissertation has focused on the health effects of exposure to $PM_{2.5}$ however this is only one aspect of air pollution. Other pollutants such as nitrogen dioxide or ozone have been measured by stationary AIRSpeck sensors in Delhi and the effects of exposure to these compounds could also be investigated.

While PCMCI has tremendous power and potential it is just one example of a way to use big data to carry out causal inference. Other examples already exist in literature and may be worth exploring and implementing. [11] uses machine learning methods to carry out targeted maximum likelihood estimation in order to quantify the effects of air pollution (specifically traffic-related) on instances of low birth weight. This closely relates to the use of ensemble learning and g-computation. It is also now possible to use Monte Carlo methods to compute posterior distributions - this too can be applied when examining complex systems for causal relationships.

As an aside, [47] argues that air pollution epidemiology studies based on observational data are only providing half the solution: it goes on to suggest that these studies should be augmented by investigating the effects of well-defined regulatory interventions in order to get a clearer picture of the potential outcomes of a particular method of pollution control. Such an investigation may well naturally arise in the DAPHNE study as, at the time of writing, the ongoing COVID-19 pandemic has resulted in the effective lockdown of many countries across the globe including India. Air quality in major cities has shown a dramatic improvement - satellite imaging of China and areas of northern Italy in March 2020 indicates large reductions in nitrogen dioxide concentrations from 2019 ([51], [50]). If similar reductions have taken place in India they may be apparent in the DAPHNE data, after which the health of the DAPHNE cohort should be investigated again.

As a final point we address an ongoing research area which naturally flows from the work in this project: estimating the strength of causal relationships. Once the existence

of an exposure-response relationship has been proven using methods such as Granger Causality or PCMCI it is important to quantify its strength. Relatively simple methods of doing so have been discussed for a long time: in the case of linear dependencies we can simply carry out a multivariate regression of a response on its causal parents, interpreting the coefficients of the fit as a measure of the strength of the relationship. More sophisticated measures of information transfer are being researched: in fact in [36] it is argued that the MCI test statistic calculated in the second stage of PCMCI can be used to rank causal parents in order of strength. PCMCI itself ranks links based on the conditional independence test statistic in order to choose the strongest dependencies. Other measures of causal strength take a more information theoretic approach, for example [39] uses transfer entropy to motivate a new concept of *momentary information transfer*. Applying these methods to the DAPHNE data can not only identify which drivers of respiratory rate are strongest, they could also be used to determine the delay before ”peak” response by evaluating the strength of causal links at different lag lengths. This information could prove invaluable in the domain of response prediction.

It is clear from this dissertation that the DAPHNE study is already bearing fruit when it comes to analysing the health effects of air pollution. We look forward to future results of the project.

Appendices

Appendix A

Full Results of Parametric Granger Tests

The following is a full table of p-values obtained from the VAR-based Granger Causality tests for linear relationships. Each row shows p-values (rounded to 5.d.p.) obtained from Granger tests for a given trial at lag lengths of 1, 5, 10, 15, 30, 45 and 60 minutes. P-values are also colour-coded as follows:

- $p < 0.05$: Green.
- $0.05 \leq p < 0.1$: Yellow.
- $0.1 \leq p$: Red.

NaN values are reported where the test failed or the ADF test for stationarity failed to reject the null hypothesis of non-stationarity.

Trial	1	5	10	15	30	45	60
DAP022(2)	0.25747	0.41189	0.37503	0.73599	0.81335	0.44814	0.15908
DAP001(1)	0.06099	0.00631	0.00802	0.06219	0.28704	0.11878	0.04021
DAP080(1)	0.27281	0.32515	0.04841	0.22312	0.27424	0.63057	0.79293
DAP040(2)	0.19166	0.13616	0.00131	2E-05	0.00053	0	1E-05
DAP075(2)	0.31745	0.64051	0.65206	0.70684	0.50122	0.28922	0.34637
DAP028(1)	0.02148	0.12011	0.40312	0.24933	0.47378	0.27721	0.22068
DAP038(1)	0.88978	0.33079	0.04333	0.04791	0.08755	0.00647	0.02015
DAP087(1)	0.82909	0.83726	0.84892	0.89331	0.92951	0.85024	0.7324
DAP101(1)	0.00184	0.06058	0.1979	0.02232	0.00138	0.00202	0.00209
DAP002(1)	0.19676	0.58457	0.58736	0.63025	0.88035	0.76303	0.63288
DAP125(1)	0.24635	0.42775	0.40951	0.22599	0.75683	0.90822	0.91912
DAP089(1)	0.35753	0.06105	0.13533	0.0803	0.44845	0.66659	0.47737
DAP016(1)	0.70335	0.91024	0.32401	0.29271	0.69438	0.86877	0.97107
DAP072(1)	0.75093	0.95178	0.9405	0.93188	0.66002	0.49612	0.19624
DAP055(1)	nan	nan	nan	nan	nan	nan	nan
DAP060(1)	0.80006	0.67097	0.94552	0.83085	0.67729	0.55048	0.36951
DAP021(1)	0.64184	0.44308	0.56837	0.66362	0.67814	0.78415	0.29163
DAP025(1)	0.14071	0.68181	0.93516	0.68748	0.92057	0.93974	0.60272
DAP103(1)	0.00405	0.07058	0.1192	0.13838	0.0006	0.00012	0.00146
DAP006(3)	0.43724	0.69488	0.0624	0.17192	0.02089	0.01963	7E-05
DAP012(1)	0.85832	0.38547	0.04922	0.00522	1E-05	8E-05	0.00207
DAP072(2)	0.45917	0.3234	0.05243	0.15154	0.2324	0.16329	0.05946
DAP074(1)	0.0579	0.00185	0.01733	0.04834	0.00462	0.00556	0.02702
DAP015(1)	0.68953	0.02173	0.0262	0	0	0	0
DAP088(1)	0.9353	0.10539	0.03077	0.15917	0.25698	0.07111	0.15135
DAP078(2)	0.06363	0.70663	0.78793	0.19623	0.00096	0.0003	0
DAP102(1)	0.2606	0.15013	0.30455	0.66065	0.98838	nan	nan
DAP001(3)	0.49075	0.86895	0.98346	0.23211	0.12434	0.2817	0.02701
DAP056(1)	0.67023	0.32626	0.32024	0.53819	0.38219	0.53525	0.28449
DAP026(1)	0.68147	0.90019	0.71382	0.94903	0.36926	0.74475	0.72632
DAP071(1)	0.98867	0.15166	0.05589	0.11249	0.21312	0.15516	0.20372
DAP111(1)	0.26952	0.21406	0.27002	0.33295	0.23178	0.206	0.3179
DAP052(1)	0.00224	0.01397	0.02079	0.03134	0.00674	0.00152	0.00058
DAP127(1)	0.21297	0.06657	0.43549	0.26415	0.17019	0.21929	0.28579
DAP126(1)	0.01165	0.10518	0.14709	0.29993	0.45401	0.83098	0.93978
DAP076(2)	0.84425	0.37569	0.29226	0.23706	0.52046	0.72845	0.83286
DAP006(1)	0.50301	0.46739	0.38671	0.69306	0.00187	1E-05	7E-05
DAP068(1)	0.94778	0.13499	0.16404	0.02429	0	0	0
DAP070(1)	0.01689	9E-05	0.00041	0.00491	3E-05	5E-05	0.00085
DAP108(1)	0.23994	0.74986	0.04009	0.10981	0.03766	0.01115	0.02205
DAP117(1)	0.82926	0.73205	0.86215	0.21901	0.32532	0.34421	0.42533
DAP055(3)	0.26559	0.04829	0.0243	6E-05	1E-05	0	3E-05
DAP095(2)	0.4144	0.49707	0.64892	0.18099	0.3564	0.51121	0.26754
DAP036(1)	0.66896	0.92789	0.45465	0.72339	0.3409	0.36917	0.2959
DAP039(1)	0.18239	0.09101	0.29323	0.22932	0.41526	0.77626	0.40598
DAP084(1)	0.34313	0.60323	0.01533	0.01775	0.10187	0.00295	0.00012
DAP081(1)	0.32075	0.64408	0.94508	0.99037	0.96115	0.95036	0.97003
DAP066(1)	0.72054	0.8769	0.87401	0.31337	0.50042	0.88272	0.84039
DAP075(1)	0.36944	0.85671	0.90478	0.72604	0.02851	0.01747	0.01561
DAP085(2)	0.24231	0.14162	0.07185	0.03222	0.00107	0.00047	0.00018
DAP062(1)	0.18608	0.03526	0.30888	0.18446	0.51883	0.30481	0.62527
DAP031(2)	0.97233	0.93097	0.48831	0.36219	0.2071	0.18566	0.18545
DAP054(1)	0.22843	0.80694	0.76834	0.74948	0.67652	0.4819	0.39694
DAP069(2)	0.29069	0.11764	0.19624	0.21948	0.24703	0.01586	0.17785
DAP091(1)	0.54197	0.71359	0.21927	0.32416	0.63856	0.79906	0.65629
DAP078(1)	0.17005	0.37154	0.63948	0.63069	0.8369	0.46403	0.51197
DAP011(1)	0.97785	0.72045	0.70111	0.9451	0.93658	0.99562	0.72549
DAP086(1)	0.15749	0.04405	0.00185	0.0102	0.00036	0	0
DAP047(1)	1	1	1	1	1	1	1
DAP037(2)	0.61015	0.56893	0.81943	0.6962	0.38117	0.36789	0.43261
DAP007(2)	0.29517	0.75743	0.43896	0.11993	0.31117	0.32667	0.37468
DAP033(2)	0.28036	0.60577	0.28732	0.11056	5E-05	0	0
DAP096(2)	0.90035	0.81825	0.86004	0.8123	0.62208	0.27955	0.08617
DAP090(1)	0.876	0.12582	0.06695	0.12837	0.22783	0.38699	0.19321
DAP012(2)	0.35405	0.67223	0.13342	0.03078	0.0726	0.09286	0.22906
DAP019(1)	0.5563	0.4436	0.591	0.48777	0.65289	0.62444	0.77308
DAP068(2)	0.63059	0	0	0	0	0	0
DAP055(2)	0.77813	0.00449	0.00211	1E-05	0.00016	8E-05	3E-05
DAP097(2)	0.43208	0.2456	0.2266	0.47944	0.08082	0.07706	0.13268
DAP115(1)	0.58709	0.88334	0.72155	0.81082	0.80225	0.9854	0.99629
DAP079(1)	0.29585	0.36981	0.00228	0.01469	0.0001	0.00015	0
DAP048(2)	0.58587	0.58536	0.93884	0.80145	0.4074	0.20609	0.16218
DAP109(1)	0.17654	0.03549	0.06626	0.25158	0.18163	0.23601	0.17205
DAP116(1)	0.4689	0.0296	0.07154	0.08398	0.00305	0.00425	0.00296
DAP008(1)	0.87716	0.00054	3E-05	0	0	0	0
DAP120(1)	0.01484	0.01512	0.0015	0.00264	0.0249	0.0359	0.01636
DAP029(1)	0.05795	0.54192	0.02818	0.06632	0.04851	0.02755	0.27904
DAP022(1)	0.41976	0.44527	0.55962	0.08611	0.08576	0.14039	0.24483
DAP031(1)	0.02171	7E-05	0.00458	0.00869	0.03648	0.0241	0.01013
DAP010(2)	0.00679	0.07879	0.03462	0.16113	0.02624	0.00488	0.02179
DAP080(2)	0.25007	0.55505	0.63521	0.55067	0.01863	0.00216	0.00033
DAP050(1)	0.08451	0.05731	0.23239	0.49871	0.31166	0.39122	0.1706
DAP043(1)	0.98554	0.49102	0.72214	0.71662	0.66253	0.26911	0.53981
DAP056(2)	0.16406	0.35497	0.18322	0.17749	0.17752	0.2607	0.62351
DAP022(3)	0.86703	0.0616	0.01427	0.02313	0.03066	0.00178	0.0004
DAP044(1)	0.51278	0.89137	0.87393	0.6067	0.17014	0.58588	0
DAP070(2)	0.00032	0.00489	0.00188	0.00101	0.00042	0.00209	1E-05
DAP050(2)	0.54824	0.89662	0.99806	0.58639	0.09605	0.00432	0.00865
DAP112(1)	0.95312	0.12639	0.31993	0.65078	0.58561	0.1322	0.18535
DAP048(1)	0.14135	0.00297	0.00146	0.00182	5E-05	0	0
DAP007(3)	0.90013	0.56788	0.19031	0.14501	0.01344	0.00021	0
DAP064(1)	0.93056	0.06698	0.26198	0.31072	0.06473	0.00498	1E-05

DAP057(3)	0.13615	0.20765	0.71999	0.28789	0.49573	0.54564	0.6706
DAP047(2)	0.09042	0.38379	0.07809	0.0229	0.00701	0.06314	0.00656
DAP037(1)	0.7868	0.66667	0.44788	0.28433	0.22203	0.19075	0.30518
DAP095(1)	0.76628	0.78511	0.82056	0.83657	0.93651	0.94498	0.97668
DAP008(3)	0.70012	0.79186	0.87168	0.72285	0.92344	0.77122	0.8528
DAP020(1)	0.99106	0.93566	0.93624	0.9109	0.99977	1	1
DAP041(1)	0.33817	0.02431	0.02666	0.01187	0.00367	0.14485	0.16428
DAP118(1)	0.14003	0.24882	0.01823	0.00212	0	0	0
DAP067(1)	0.98808	0.54533	0.8058	0.69439	0.32239	0.57207	0.3218
DAP063(1)	0.22561	0.7666	0.15088	0.18867	0.2091	0.22493	0.47431
DAP057(1)	0.54539	0.62202	0.42947	0.22071	0.63083	0.66191	0.44424
DAP014(2)	0.03195	0.05167	0.00365	0.00184	0.00038	0.00191	0.01567
DAP073(2)	0.00461	0.02955	0.02493	0.00804	0.0059	0.03801	0.03343
DAP032(1)	0.92615	0.91248	0.33972	0.73061	0.55718	0.52736	0.84082
DAP123(1)	0.87341	0.6004	0.89962	0.84729	0.95046	0.8067	0.61413
DAP122(1)	0.63174	0.14051	0.09032	0.11678	0.46467	0.32339	0.0752
DAP067(2)	0.35929	0.1416	0.14737	0.32431	0.54084	0.62212	0.71693
DAP076(1)	0.03005	0.00826	9E-05	0	0	0	0
DAP119(1)	0.30861	0.69191	0.8588	0.71743	0.43989	0.63896	0.82618
DAP005(1)	0.07623	0.15793	0.21005	0.21165	0.00345	0.00572	0.00011
DAP121(1)	0.29697	0.16316	0.08794	0.18621	0.01407	0.02843	0.17526
DAP007(1)	0.24747	0.68936	0.49197	0.62971	0.68123	0.92235	0.59504
DAP114(1)	0.82079	0.99814	0.99973	0.76044	0.71876	0.80994	0.37486
DAP004(1)	0.40344	0.11371	0.27075	0.31323	0.655	0.76097	0.97971
DAP043(2)	0.47938	0.35047	0.64562	0.31318	0.60012	0.11115	0.12233
DAP018(1)	0.90282	0.445	0.80214	0.74604	0.31	0.53979	0.34648
DAP024(1)	0.3406	0.03638	0.00582	4E-05	0	0	0
DAP033(1)	0.66318	0.358	0.08343	0.06249	0.40436	0.90098	0.9811
DAP018(2)	0.56277	0.63064	0.04071	0.0002	0.00107	2E-05	2E-05
DAP094(1)	0.08725	0.0658	0.06612	0.22429	0.47862	0.44258	0.67587
DAP042(2)	0.81269	0.98783	0.77697	0.75588	0.23845	0.3676	0.60588
DAP065(1)	0.00705	0.00166	0.00653	0	nan	nan	nan
DAP098(2)	0.57148	0.12708	0.35985	0.63834	0.56644	0.05789	0.11768
DAP028(2)	0.87855	0.91129	0.68701	0.9388	0.97384	0.90064	0.92893
DAP009(1)	0.28601	0.11116	0.32884	0.67384	0.2376	0.09786	0.16164
DAP069(1)	0.38563	0.84232	0.27267	0.52295	0.15271	0.33695	0.16931
DAP061(1)	0.02429	0.00891	0.01632	0.0728	0.2566	0.37838	0.31465
DAP010(1)	0.27756	0.01082	0.00253	0.0003	8E-05	2E-05	0
DAP034(1)	0.17743	0.91327	0.62675	0.50504	0.19186	0.08916	0.03211
DAP013(2)	0.00221	0	0	0	0	0	0
DAP042(1)	0.02868	0.21071	0.15326	0.41629	0.67541	0.74514	0.75978
DAP036(2)	0.5617	0.12375	0.07921	0.02209	0.01698	0.02427	0.00181
DAP011(2)	0.04642	0.25366	0.01584	0.00043	0.0001	0.00019	0.00373
DAP092(1)	0.15049	0.55419	0.40761	0.52239	0.10341	0.12807	0.19758
DAP084(2)	0.83148	0.2995	0.45175	0.3004	0.58312	0.94016	0.86173
DAP113(1)	0.67519	0.16956	0.35936	0.4882	0.32635	0.16393	0.03088
DAP046(1)	0.83097	0.67349	0.12404	0.00816	0.00025	0.00024	1E-05
DAP105(1)	0.14319	0.84581	0.8698	0.9271	0.60034	0.87704	0.75954
DAP011(3)	0.82805	0.83089	0.8346	0.77817	0.92711	0.95689	0.8754
DAP051(1)	0.56544	0.59467	0.6036	0.88564	0.71012	0.6725	0.5132
DAP057(2)	0.27917	0.93686	0.90703	0.86708	0.8049	0.53763	0.45489
DAP008(2)	0	0	0	0	0	0	0
DAP110(1)	0.89594	0.14495	0.23629	0.23679	0.32839	0.15322	0.12019
DAP004(2)	0.2228	0.57723	0.84335	0.37811	0.04181	0.00033	0.00449
DAP023(1)	0.14394	0.33475	0.43515	0.00604	0.09514	0.17899	0.23144
DAP077(1)	0.501	0.29043	0.492	0.1101	0.11488	0.24096	0.32564
DAP001(2)	0.86726	0.57083	0.58214	0.05523	0.09759	0.00657	0.00125
DAP003(2)	0.76461	0.83218	0.1386	0.0004	0.00011	0	0
DAP049(1)	0.22544	0.73869	0.03626	0.00653	0.00063	1E-05	2E-05
DAP053(1)	0.92904	0.99718	0.68204	0.86279	0.8474	0.86418	0.80439
DAP058(1)	0.84074	0.7826	0.29421	0.30577	0.58342	0.82498	0.73077
DAP106(1)	0.98725	0.19797	0.41321	0.32246	0.61586	0.46372	0.59098
DAP083(1)	0.37026	0.31636	0.01742	0.00302	0.00043	5E-05	0.00013
DAP006(2)	0.22263	0.2297	0.43924	0.03748	0.04645	0.02792	0.00401
DAP056(3)	0.58189	0.16379	0.39333	0.21774	0.2237	0.44598	0.30305
DAP013(1)	0.26057	0.3985	0.87736	0.74195	0.73851	0.71774	0.87445
DAP002(2)	nan	nan	nan	nan	nan	nan	nan
DAP035(1)	0.09297	0.42926	0.29519	0.05484	0.07076	0.04942	0.00794
DAP027(1)	0.17799	0.30611	0.16748	0.09063	0.21954	0.36901	0.55595
DAP005(2)	0.71745	0.73467	0.11923	0.01549	0.00281	0.00411	0.0001
DAP045(1)	0.63835	0.39574	0.25319	0.27718	0.15412	0.05313	0.23851
DAP017(1)	0.77626	0.65048	0.25539	0.16743	0.08293	0.10611	0.11225
DAP032(2)	0.26098	0.03964	0.0903	0.29586	0.18368	0.15519	0.05794
DAP085(1)	0.3519	0.2787	0.83211	0.84715	0.84645	0.07067	0.04668
DAP030(1)	0.57789	0.44381	0.49696	0.69306	0.47582	0.43696	0.29784
DAP093(1)	0.14576	0.3735	0.17229	0.30769	0.22997	0.25554	0.36411
DAP059(2)	0.52534	0.68737	0.35011	0.28422	0.14486	0.38754	0.67655
DAP096(1)	0.70325	0.50725	0.15356	0.05583	0.00029	0	0
DAP086(2)	0.33882	0.6852	0.72904	0.56941	0.70288	0.46458	0.69119
DAP124(1)	0.32467	0.52939	0.69829	0.23305	0.20797	0.11621	0.02335
DAP089(2)	0.99385	0.85293	0.99421	0.88564	0.12207	0.00803	0.00159
DAP107(1)	0.81291	0.70615	0.45019	0.81695	0.96588	0.86312	0.39663
DAP014(1)	0.39956	0.98142	0.93529	0.97662	0.78274	0.90257	0.5519
DAP003(1)	0.93073	0	0	0	0	0	0
DAP059(1)	0.13147	0.50233	0.88691	0.71835	0.68211	0.63809	0.53731
DAP074(2)	0.1438	0.00818	0.00017	1E-05	0	0	0
DAP021(2)	0.35189	0.05877	0.1326	0.11426	0.09558	0.01823	0.00122
DAP017(2)	0.59095	0.73746	0.92336	0.84788	0.79505	0.94313	0.48784
DAP104(1)	0.38431	0.50987	0.20111	0.10347	0.25477	0.42004	0.27871
DAP082(1)	0.3545	0.8986	0.9469	0.5808	0.61751	0.3734	0.5339
DAP040(1)	0	0	0	0	0	0	0

Appendix B

Full Results of Non-Parametric Granger Tests

The following is a full table of p-values obtained from the Diks-Panchenko tests for non-linear Granger Causality. Each row shows p-values (rounded to 5.d.p.) obtained from Granger tests for a given trial at lag lengths of 1, 5, 10, 15, 30, 45 and 60 minutes. P-values are also colour-coded as follows:

- $p < 0.05$: Green.
- $0.05 \leq p < 0.1$: Yellow.
- $0.1 \leq p$: Red.

NaN values are reported where the test failed.

Trial	1	5	10	15	30	45	60
DAP022(2)	0.99939	0.99936	0.99584	0.99158	0.82087	0.8349	0.87342
DAP001(1)	0.00177	3E-05	0.00029	0.00184	0.24891	0.42055	0.32208
DAP080(1)	0.35176	0.31558	0.16314	0.08526	0.16986	0.36783	0.62317
DAP040(2)	0.31036	0.18137	0.07563	0.156	0.15085	0.4065	0.59488
DAP075(2)	0.59822	0.45254	0.47276	0.50455	0.35401	0.46047	0.57993
DAP028(1)	0.33321	0.41019	0.38224	0.29529	0.06187	0.07146	0.12832
DAP038(1)	0.22644	0.241	0.18829	0.01933	0.04013	0.21274	0.45558
DAP087(1)	0.33883	0.02113	0.0008	0.00016	0.00032	0.00404	0.02236
DAP101(1)	0.46155	0.00552	0.00053	0.00198	0.0327	0.08083	0.02849
DAP002(1)	0.6133	0.15138	0.37617	0.64133	0.1654	0.15279	0.29918
DAP125(1)	0.68271	0.07376	0.00812	0.01208	0.16969	0.19563	0.40114
DAP089(1)	0.8904	0.93825	0.87149	0.53962	0.04549	0.0611	0.04064
DAP016(1)	0.13517	0.07464	0.18208	0.2401	0.51076	0.85845	0.86125
DAP072(1)	0.06523	0.03735	0.05795	0.05139	0.03323	0.00759	0.00617
DAP055(1)	0.84931	0.89694	0.71375	0.81021	0.47344	0.6259	0.5981
DAP060(1)	0.07821	0.08345	0.0099	0.0064	0.00235	0.01478	0.12325
DAP021(1)	0.61856	0.08803	0.0515	0.04029	0.04997	0.08621	0.06494
DAP025(1)	0.67994	0.82244	0.89154	0.83721	0.58293	0.23411	0.34747
DAP103(1)	0.07227	0.00564	0.00568	0.16097	0.01013	0.01231	0.1111
DAP006(3)	0.28654	0.04388	0.00439	0.00254	0.00194	0.00057	0.00438
DAP012(1)	0.38561	0.0033	0.00035	0.00097	0.00023	0.00036	0.00323
DAP072(2)	0.03339	0.00107	0.00095	0.00314	0.00015	0.00016	0.00065
DAP074(1)	0.00097	0	0	2E-05	0.01094	0.02226	0.00015
DAP015(1)	0.00031	0	0	0	0	0	6E-05
DAP088(1)	0.57723	0.1656	0.05203	0.4041	0.35633	0.32131	0.37345
DAP078(2)	0.01674	0.00267	0.02854	0.04112	0.23792	0.18155	0.1537
DAP102(1)	0.66854	0.44468	0.02298	0.15391	0.23392	nan	nan
DAP001(3)	0.39879	0.15934	0.08497	0.03553	0.08566	0.30723	0.04751
DAP056(1)	0.42739	0.07323	0.23058	0.57415	0.75243	0.50117	0.3109
DAP026(1)	0.43953	0.7055	0.86504	0.95168	0.21496	0.11226	0.69375
DAP071(1)	0.00785	0.0001	0.0001	0.00012	8E-05	0.00011	0.00088
DAP111(1)	0.79645	0.13444	0.16206	0.0386	0.00864	0.16868	0.04844
DAP052(1)	0.02941	0.02011	0.03701	0.10028	0.05204	0.09409	0.06034
DAP127(1)	0.35807	0.09576	0.06846	0.05341	0.27535	0.22597	0.28465
DAP126(1)	0.94987	0.99335	0.97678	0.74438	0.21675	0.12246	0.29738
DAP076(2)	0.27956	0.15194	0.23731	0.4056	0.68756	0.93751	0.94493
DAP006(1)	0.10174	0.0342	0.07091	0.12055	0.02784	0.00493	0.00504
DAP068(1)	0.02696	0.00195	0.00526	0.00963	0.04946	0.06645	0.05056
DAP070(1)	0.03932	0.01062	0.00053	0.00027	0.00206	0.037	0.02565
DAP108(1)	0.19259	0.515	0.41176	0.14814	0.16713	0.01631	0.00631
DAP117(1)	0.59943	0.16291	0.24851	0.25635	0.27626	0.42427	0.16629
DAP055(3)	0.05403	0.00639	0.00533	0.00222	0.02618	0.08978	0.11523
DAP095(2)	0.45412	0.0797	0.14671	0.04572	0.0087	0.02625	0.1768
DAP036(1)	0.25249	0.21625	0.04817	0.02879	0.02102	0.00756	0.02072
DAP039(1)	0.11304	0.07703	0.0958	0.28344	0.10097	0.22482	0.23439
DAP084(1)	0.0223	7E-05	0	0	0.00015	0.01117	0.21362
DAP081(1)	0.84149	0.90478	0.35119	0.22672	0.15086	0.03106	0.09858
DAP066(1)	0.0017	0.00055	0.00461	0.00704	0.03897	0.04354	0.1252
DAP075(1)	0.62892	0.70432	0.75012	0.43946	0.4754	0.65483	0.7993
DAP085(2)	0.03899	0.00992	0.00147	0.00326	0.02415	0.19899	0.33871
DAP062(1)	0.02448	0.00203	0.00115	0.00425	0.00354	0.00331	0.00241
DAP031(2)	0.19788	0.10777	0.04137	0.01971	0.022	0.03047	0.06739
DAP054(1)	0.31065	0.16679	0.29398	0.0982	0.00338	0.04321	0.02839
DAP069(2)	0.33423	0.12057	0.45981	0.32189	0.27284	0.7406	nan
DAP091(1)	0.00329	0.00459	0.00483	0.00573	0.11295	0.16706	0.54536
DAP078(1)	0.9651	0.97146	0.99888	0.99932	0.99953	0.9997	0.99642
DAP011(1)	0.2758	0.53341	0.25925	0.45299	0.4926	0.41303	0.3002
DAP086(1)	0.02475	0.0009	0.00689	0.11097	0.08001	0.04864	0.02002
DAP047(1)	1	1	1	0	0	3E-05	0.00063
DAP037(2)	0.01689	0.00747	0.05693	0.18534	0.03936	0.02171	0.03893
DAP007(2)	0.69163	0.84929	0.87267	0.92157	0.65376	0.48011	0.65448
DAP033(2)	9E-05	0	0	0	0	0	0
DAP096(2)	0.13019	0.03412	0.00326	0.00476	0.02738	0.15748	0.36408
DAP090(1)	0.0307	0.05961	0.03307	0.00673	0	0	1E-05
DAP012(2)	0.00875	2E-05	3E-05	3E-05	0.00891	0.01057	0.09996
DAP019(1)	0.0996	0.06811	0.03465	0.03978	0.25555	0.28431	0.36812
DAP068(2)	0.00753	0.08796	0.09786	0.11118	0.07882	0.08663	0.03821
DAP055(2)	0.01428	0.02673	0.01226	0.00397	0.01464	0.0401	0.10646
DAP097(2)	0.67225	0.29674	0.0133	0.00506	0.00201	0.00234	0.01481
DAP115(1)	0.47574	0.74757	0.42664	0.65613	0.67888	0.75592	0.78999
DAP079(1)	0.00051	0	0.00379	0.02326	0.01634	0.00743	0.01265
DAP048(2)	0.93121	0.45614	0.29086	0.13662	0.13993	0.14192	0.07789
DAP109(1)	0.44267	0.55784	0.72691	0.62619	0.69285	0.53064	0.87005
DAP116(1)	0.01418	0.00028	0.00031	0.0012	0.00455	0.00286	0.0037
DAP008(1)	0.13272	0.16019	0.03847	0.03969	0.04047	0.07824	0.08062
DAP120(1)	0.01415	0.00501	0.00258	0.00025	0.00345	0.00032	0.00293
DAP029(1)	0.16635	0.04569	0.01943	0.04277	0.26946	0.36328	0.40371
DAP022(1)	0.93337	0.12986	0.04484	0.32024	0.87927	0.74182	0.87708
DAP031(1)	0.2021	0.06062	0.08706	0.04733	0.03495	0.02779	0.05196
DAP010(2)	0.09826	0.19259	0.38854	0.40457	0.15271	0.0222	0.12757
DAP080(2)	0.12394	0.08125	0.3577	0.12442	0.25828	0.04121	0.03312
DAP050(1)	0.03851	0.03143	0.02402	0.00809	0.01562	0.03138	0.10034
DAP043(1)	0.09838	0.04492	0.01671	0.01536	0.00451	0.00889	0.00901
DAP056(2)	0.08892	0.02311	0.10101	0.06127	0.01158	0.0853	0.38814
DAP022(3)	0.00013	0	0	3E-05	0.01916	0.01251	0.00258
DAP044(1)	0.86054	0.81585	0.96038	0.8458	0.9529	0.14319	0.54183
DAP070(2)	0.01144	0.00272	0.00054	0.00049	0.0083	0.0083	0.05881
DAP050(2)	0.0588	0.08984	0.00173	0.00133	0.01201	0.01568	0.11416
DAP112(1)	0.22249	0.2313	0.20982	0.14106	0.34066	0.55501	0.55968
DAP048(1)	0.12501	0.10769	0.06206	0.00968	0.04051	0.02681	0.05675
DAP007(3)	0.05599	0.00052	6E-05	0.0007	0.00702	0.00761	0.02466
DAP064(1)	0.39008	0.00326	0.00055	0.00237	0.00173	0.0018	0.00299

DAP057(3)	0.32462	0.45508	0.08878	0.10806	0.34386	0.3202	0.78218
DAP047(2)	0.04602	0.00913	0.00036	0.00032	0.01593	0.00117	0.00387
DAP037(1)	0.18061	0.25887	0.24973	0.22719	0.19903	0.21005	0.08918
DAP095(1)	0.55043	0.52098	0.25697	0.0375	0.1514	0.23722	0.84193
DAP008(3)	0.41273	0.65162	0.7326	0.61335	0.72949	0.67871	0.77977
DAP020(1)	0.95116	0.91424	0.54281	0.7538	0.97396	0.99601	0.99713
DAP041(1)	0.14971	0.01081	0.00312	0.02772	0.22261	0.09536	0.00892
DAP118(1)	3E-05	0	3E-05	8E-05	0.00642	0.04185	0.01793
DAP067(1)	0.01575	0.01371	0.00453	0.00371	0.02552	0.15638	0.27125
DAP063(1)	0.07885	0.06762	0.00823	0.00238	0.00138	0.00022	0.0006
DAP057(1)	0.31302	0.40162	0.2194	0.23463	0.1423	0.14248	0.30034
DAP014(2)	0.09476	0.5602	0.16768	0.14892	0.04226	0.19474	0.20957
DAP073(2)	0.02585	0.00012	0	0	0	7E-05	0.00061
DAP032(1)	0.68565	0.53327	0.43427	0.46447	0.34676	0.40877	0.78133
DAP123(1)	0.10902	0.26761	0.06084	0.0221	0.04173	0.02275	0.08861
DAP122(1)	0.00032	0	0	0	0.01406	0.05059	0.01678
DAP067(2)	0.06987	0.0219	0.01165	0.00416	0.00772	0.01433	0.01413
DAP076(1)	0.14364	0.03887	0.01053	0.01153	0.01163	0.01105	0.02047
DAP119(1)	0.02862	0.00063	0.0009	0.00344	0.00473	0.0327	0.00734
DAP005(1)	0.08289	0.04663	0.0737	0.07364	0.1972	0.08822	0.05916
DAP121(1)	0.1186	0.00319	0.00268	0.00892	0.02123	0.06444	0.11641
DAP007(1)	0.06731	0.0046	0.0176	0.08408	0.21885	0.45427	0.9001
DAP114(1)	0.52022	0.00414	0	2E-05	0.00039	0.00059	0.00105
DAP004(1)	0.11567	0.30946	0.39291	0.23191	0.1135	0.35748	0.69399
DAP043(2)	0.43256	0.57279	0.40492	0.38978	0.08848	0.73331	0.07011
DAP018(1)	0.73655	0.48894	0.50835	0.49931	0.4904	0.17572	0.2508
DAP024(1)	0.00209	0.00059	0.01882	0.11018	0.05301	0.06375	0.05679
DAP033(1)	0.41057	0.18683	0.02048	0.01863	0.1732	0.15293	0.20009
DAP018(2)	0.04638	0.05244	0.00622	0.00072	0.00649	0.00749	0.02812
DAP094(1)	0.59503	0.73695	0.49103	0.44727	0.69354	0.14526	0.12675
DAP042(2)	0.10077	0.03205	0.16405	0.31901	0.04004	0.0872	0.04947
DAP065(1)	0.25783	0.1141	0.04298	0.00922	0.0535	0.22182	0.23816
DAP098(2)	0.96652	0.18493	0.17568	0.06922	0.02045	0.02071	0.06765
DAP028(2)	0.64401	0.38367	0.57771	0.55409	0.07156	0.1463	0.50848
DAP009(1)	0.11472	0.09556	0.41639	0.43496	0.44952	0.66672	0.89774
DAP069(1)	0.03609	0.00681	0.00058	8E-05	0.00688	0.01029	0.00838
DAP061(1)	0.03108	0.00217	0.0003	1E-05	6E-05	0.00016	0.00061
DAP010(1)	0.74008	0.00879	0.02032	0.0159	0.04779	0.01529	0.01099
DAP034(1)	0.18484	0.08497	0.02293	0.01602	0.00587	0.01068	0.0126
DAP013(2)	0.0625	0.0033	0.00117	0.00092	0.00171	0.00363	0.01403
DAP042(1)	0.71883	0.76803	0.70071	0.68608	0.56683	0.78017	0.41894
DAP036(2)	0.00175	0	0	0	9E-05	0.00028	0.00128
DAP011(2)	0.09625	0.03811	0.04433	0.03563	0.07268	0.01893	0.01407
DAP092(1)	0.06075	0.1526	0.15971	0.17669	0.51361	0.39272	0.70867
DAP084(2)	0.49128	0.21994	0.12521	0.00328	0.00712	0.04591	0.28403
DAP113(1)	0.04785	0.0063	0.00528	0.00206	0.00595	0.04488	0.0436
DAP046(1)	0.00041	0	0	0	0	0	0.00019
DAP105(1)	0.041	0.03507	0.07126	0.07736	0.11117	0.40046	0.73613
DAP011(3)	0.11914	0.12758	0.14135	0.07492	0.11504	0.02367	0.09336
DAP051(1)	0.92355	0.07222	0.00407	0.00093	0.00768	0.00047	0.00037
DAP057(2)	0.49792	0.26489	0.47213	0.67657	0.42146	0.48847	0.23152
DAP008(2)	0.04298	0.01294	0.00415	0.00232	0.00114	0.00264	0.00359
DAP110(1)	0.17401	0.00493	0.00543	0.00063	0.00372	0.00522	0.02026
DAP004(2)	0.0898	0.0661	0.03627	0.0572	0.07222	0.09393	0.17104
DAP023(1)	0.16916	0.60439	0.56808	0.4452	0.52248	0.68221	0.26676
DAP077(1)	0.02557	0.01888	0.03756	0.05611	0.07063	0.30008	0.54977
DAP001(2)	0.4238	0.08812	0.07643	0.12323	0.37385	0.69257	0.77932
DAP003(2)	0.09047	0.02965	0.06632	0.075	0.08356	0.17856	0.08006
DAP049(1)	0.88521	0.6017	0.22611	0.15984	0.03463	0.07959	0.1201
DAP053(1)	0.22011	0.12952	0.1104	0.195	0.58953	0.6578	0.26665
DAP058(1)	0.76793	0.35286	0.2568	0.04614	0.01414	0.16677	0.24828
DAP106(1)	0.10791	0.03131	0.01026	0.01307	0.13762	0.42779	0.35906
DAP083(1)	0.00012	0	0	0	0	0	6E-05
DAP006(2)	0.02422	0.01177	0.03978	0.08095	0.21263	0.32607	0.08131
DAP056(3)	0.99601	0.33075	0.49692	0.52188	0.04219	0.05497	0.1463
DAP013(1)	0.97777	0.80064	0.80183	0.86536	0.85853	0.93685	0.94918
DAP002(2)	1	0	1	1	1	0.00443	0.28288
DAP035(1)	0.10311	0.09456	0.04329	0.03466	0.23294	0.70612	0.79427
DAP027(1)	0.5503	0.87838	0.64752	0.62991	0.98581	0.99371	0.99625
DAP005(2)	0.14179	0.08072	0.02718	0.00912	0.02016	0.02609	0.12096
DAP045(1)	0.88023	0.15758	0.05533	0.00566	0.01291	0.0401	0.07086
DAP017(1)	0.30198	0.55494	0.26375	0.05659	0.16655	0.2125	0.24242
DAP032(2)	0.03268	0	0	0	0.0001	0.00026	0.00142
DAP085(1)	0.02643	0.03062	0.02528	0.00722	0.00426	0.0015	0.00384
DAP030(1)	0.2125	0.75459	0.46541	0.65317	0.3723	0.78769	0.55391
DAP093(1)	0.35789	0.07417	0.00526	0.00028	0.00093	0.00166	0.00122
DAP059(2)	0.01538	0.01327	0	0	7E-05	0.00389	0.00869
DAP096(1)	0.54336	0.22331	0.12451	0.04292	0.03375	0.05603	0.05516
DAP086(2)	0.04452	0.00011	9E-05	3E-05	0.0033	0.01369	0.02362
DAP124(1)	0.05672	0.00314	5E-05	1E-05	2E-05	0.00242	0.00102
DAP089(2)	0.04456	0.00017	0.0001	0.00145	0.05241	0.06323	0.11321
DAP107(1)	0.06152	0.0017	0.00046	0.0022	0.01401	0.01534	0.02048
DAP014(1)	0.83858	0.84005	0.94628	0.85518	0.30558	0.16418	0.3019
DAP003(1)	0.85558	0.16018	0.16001	0.08573	0.07908	0.08422	0.08227
DAP059(1)	0.8097	0.11272	0.19874	0.27334	0.05282	0.09185	0.06209
DAP074(2)	0.00241	0.00026	3E-05	0.01176	0.04702	0.09316	0.1301
DAP021(2)	0.10203	0.16145	0.23371	0.1087	0.02733	0.02139	0.01972
DAP017(2)	0.15148	0.00028	9E-05	0.0003	0.008	0.02623	0.22264
DAP104(1)	0.06401	0.01751	0.01374	0.00609	0.00601	0.00873	0.01254
DAP082(1)	0.25605	0.31194	0.31076	0.10722	0.2077	0.12323	0.19271
DAP040(1)	0.05403	0.0264	0.14003	0.45948	0.88904	0.87171	0.92424

Appendix C

Full Results of PCMCI for Linear Dependencies

The following is a full table of p-values obtained from running PCMCI with a partial correlation test for linear dependencies. Each row shows p-values (rounded to 3.d.p.) obtained from tests at all lag lengths up to 60 minutes for a given trial. P-values are also colour-coded as follows:

- $p < 0.05$: Green.
- $0.05 \leq p < 0.1$: Yellow.
- $0.1 \leq p$: Red.

Trial	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30
DAP022(2)	0.571	0.705	0.37	0.695	0.263	0.455	0.114	0.228	0.847	0.864	0.548	0.601	0.049	0.352	0.832	0.732	0.307	0.97	0.224	0.953	0.964	0.579	0.769	0.969	0.459	0.806	0.942	0.473	0.911	0.232
DAP080(1)	0.314	0.633	0.249	0.068	0.267	0.385	0.224	0.809	0.967	0.996	0.105	0.468	0.522	0.593	0.696	0.46	0.777	0.086	0.304	0.293	0.891	0.671	0.861	0.741	0.563	0.763	0.646	0.979	0.657	0.924
DAP075(2)	0.855	0.713	0.335	0.379	0.576	0.156	0.618	0.56	0.677	0.656	0.779	0.156	0.081	0.725	0.142	0.716	0.83	0.724	0.683	0.699	0.127	0.792	0.833	0.204	0.85	0.049	0.971	0.661	0.275	0.748
DAP087(1)	0.652	0.327	0.722	0.789	0.734	0.821	0.457	0.631	0.328	0.502	0.64	0.126	0.881	0.526	0.24	0.236	0.94	0.768	0.785	0.626	0.592	0.848	0.346	0.553	0.738	0.145	0.473	0.549	0.476	0.768
DAP101(1)	0.01	0.304	0.635	0.999	0.958	0.312	0.511	0.605	0.715	0.573	0.225	0.16	0.101	0.916	0.512	0.071	0.938	0.058	0.298	0.328	0.381	0.005	0.996	0.485	0.256	0.313	0.527	0.487	0.617	0.743
DAP125(1)	0.374	0.889	0.409	0.887	0.588	0.015	0.74	0.996	0.603	0.686	0.184	0.232	0.183	0.849	0.23	0.582	0.107	0.023	0.076	0.17	0.078	0.606	0.726	0.831	0.775	0.987	0.018	0.797	0.167	0.301
DAP016(1)	0.914	0.094	0.13	0.084	0.345	0.731	0.015	0.472	0.171	0.619	0.993	0.822	0.778	0.623	0.186	0.684	0.463	0.307	0.823	0.966	0.465	0.228	0.503	0.998	0.375	0.736	0.433	0.106	0.766	0.742
DAP072(1)	0.546	0.55	0.414	0.488	0.553	0.39	0.573	0.681	0.644	0.821	0.068	0.673	0.198	0.499	0.293	0.454	0.511	0.02	0.876	0.992	0.466	0.306	0.981	0.033	0.354	0.407	0.563	0.218	0.702	0.973
DAP060(1)	0.203	0.506	0.051	0.662	0.654	0.408	0.761	0.375	0.181	0.32	0.915	0.382	0.991	0.752	0.964	0.315	0.985	0.364	0.427	0.008	0.485	0.523	0.503	0.951	0.338	0.546	0.526	0.476	0.908	0.207
DAP088(1)	0.528	0.301	0.368	0.437	0.986	0.892	0.728	0.011	0.954	0.21	0.19	0.733	0.406	0.682	0.765	0.3	0.916	0.413	0.506	0.299	0.873	0.807	0.879	0.231	0.74	0.536	0.289	0.74	0.891	0.451
DAP102(1)	0.601	0.879	0.096	0.98	0.56	0.557	0.421	0.455	0.453	0.4	0.521	0.823	0.93	0.206	0.14	0.761	0.232	0.035	0.169	0.043	0.37	0.487	0.55	0.95	0.582	0.693	0.069	0.738	0.012	0.784
DAP001(3)	0.069	0.352	0.861	0.866	0.869	0.668	0.384	0.759	0.575	0.858	0.372	0.885	0	0.977	0.583	0.285	0.088	0.461	0.844	0.676	0.979	0.688	0.542	0.244	0.839	0.444	0.815	0.053	0.712	0.236
DAP126(1)	0.231	0.84	0.578	0.585	0.899	0.667	0.182	0.239	0.162	0.857	0.82	0.196	0.155	0.954	0.045	0.129	0.277	0.83	0.623	0.238	0.122	0.402	0.05	0.095	0.113	0.851	0.446	0.022	0.766	0.088
DAP108(1)	0.033	0.442	0.498	0.659	0.212	0.685	0.877	0.54	0.171	0.228	0.711	0.977	0.643	0.378	0.755	0.295	0.147	0.713	0.331	0.268	0.457	0.654	0.103	0.368	0.387	0.446	0.092	0.427	0.962	0.796
DAP095(2)	0.09	0.437	0.85	0.841	0.128	0.556	0.918	0.288	0.248	0.131	0.936	0.279	0.119	0.026	0.247	0.582	0.906	0.418	0.552	0.463	0.07	0.355	0.127	0.071	0.194	0.532	0.795	0.601	0.688	0.724
DAP084(1)	0.568	0.75	0.226	0.791	0.76	0.34	0.012	0.255	0.522	0.672	0.565	0.232	0.268	0.592	0.794	0.609	0.736	0.799	0.039	0.722	0.993	0.056	0.46	0.575	0.789	0.234	0.463	0.409	0.328	0.176
DAP081(1)	0.863	0.551	0.094	0.059	0.955	0.61	0.897	0.362	0.989	0.088	0.253	0.85	0.631	0.101	0.756	0.564	0.948	0.261	0.538	0.82	0.57	0.981	0.393	0.624	0.277	0.495	0.265	0.403	0.617	0.066
DAP075(1)	0.946	0.275	0.889	0.163	0.877	0.883	0.783	0.781	0.73	0.25	0.355	0.634	0.041	0.234	0.639	0.121	0.065	0.333	0.701	0.765	0.207	0.041	0.921	0.678	0.415	0.364	0.881	1	0.107	0.781
DAP031(2)	0.753	0.6	0.16	0.569	0.798	0.841	0.638	0.156	0.98	0.931	0.527	0.788	0.787	0.68	0.921	0.149	0.408	0.509	0.894	0.991	0.182	0.3	0.413	0.739	0.287	0.291	0.286	0.231	0.082	0.753
DAP054(1)	0.068	0.542	0.847	0.703	0.754	0.354	0.616	0.902	0.608	0.835	0.414	0.109	0.032	0.018	0.243	0.47	0.761	0.122	0.392	0.942	0.073	0.772	0.497	0.87	0.39	0.26	0.174	0.39	0.658	0.616
DAP069(2)	0.867	0.859	0.055	0.703	0.012	0.355	0.58	0.125	0.334	0.942	0.587	0.741	0.172	0.187	0.714	0.609	0.358	0.23	0.609	0.326	0.139	0.539	0.803	0.181	0.635	0.655	0.626	0.747	0.221	0.65
DAP091(1)	0.425	0.067	0.323	0.173	0.919	0.047	0.917	0.441	0.231	0.944	0.327	0.422	0.151	0.765	0.093	0.892	0.05	0.152	0.425	0.694	0.254	0.266	0.763	0.568	0.363	0.776	0.356	0.66	0.977	0.465
DAP096(2)	0.928	0.958	0.126	0.954	0.415	0.176	0.361	0.801	0.818	0.174	0.115	0.78	0.44	0.767	0.574	0.211	0.084	0.142	0.827	0.67	0.963	0.468	0.338	0.871	0.928	0.333	0.575	0.564	0.621	0.794
DAP090(1)	0.383	0.032	0.544	0.628	0.218	0.148	0.12	0.665	0.8	0.825	0.841	0.509	0.731	0.776	0.211	0.764	0.396	0.741	0.769	0.866	0.952	0.652	0.605	0.932	0.685	0.501	0.904	0.133	0.802	0.605
DAP097(2)	0.472	0.153	0.162	0.662	0.542	0.649	0.302	0.712	0.746	0.024	0.443	0.634	0.729	0.742	0.981	0.728	0.814	0.326	0.906	0.325	0.837	0.702	0.904	0.407	0.552	0.095	0.773	0.584	0.658	0.008
DAP048(2)	0.852	0.823	0.333	0.177	0.452	0.853	0.462	0.398	0.234	0.892	0.399	0.574	0.669	0.028	0.631	0.622	0.669	0.825	0.716	0.031	0.68	0.193	0.703	0.939	0.068	0.123	0.258	0.238	0.378	0.191
DAP109(1)	0.688	0.634	0.591	0.194	0.973	0.787	0.197	0.721	0.94	0.585	0.612	0.556	0.905	0.234	0.998	0.325	0.643	0.83	0.241	0.709	0.516	0.629	0.491	0.34	0.97	0.44	0.265	0.359	0.785	0.136
DAP050(1)	0.271	0.828	0.968	0.347	0.485	0.431	0.389	0.462	0.826	0.679	0.474	0.299	0.206	0.92	0.502	0.558	0.792	0.746	0.04	0.587	0.43	0.597	0.211	0.176	0.219	0.966	0.262	0.108	0.277	0.112
DAP057(3)	0.793	0.828	0.62	0.004	0.132	0.874	0.482	0.055	0.107	0.284	0.32	0.081	0.073	0.365	0.576	0.243	0.726	0.819	0.117	0.002	0.637	0.586	0.482	0.548	0.723	0.227	0.944	0.385	0.743	0.775
DAP095(1)	0.238	0.511	0.873	0.236	0.118	0.524	0.763	0.315	0.688	0.616	0.186	0.189	0.73	0.562	0.117	0.664	0.498	0.787	0.152	0.525	0.199	0.299	0.977	0.993	0.572	0.608	0.91	0.819	0.236	0.833
DAP067(1)	0.879	0.601	0.675	0.261	0.751	0.306	0.784	0.083	0.43	0.709	0.179	0.655	0.665	0.638	0.081	0.909	0.667	0.129	0.678	0.497	0.697	0.236	0.724	0.581	0.311	0.217	0.696	0.03	0.64	0.589
DAP014(2)	0.462	0.858	0.735	0.667	0.58	0.871	0.464	0.173	0.869	0.169	0.115	0.51	0.492	0.262	0.543	0.197	0.805	0.616	0.332	0.055	0.31	0.339	0.932	0.567	0.438	0.788	0.984	0.77	0.393	0.964
DAP123(1)	0.14	0.226	0.032	0.456	0.855	0.003	0.071	0.242	0.396	0.246	0.723	0.063	0.653	0.727	0.477	0.572	0.171	0.429	0.898	0.313	0.486	0.599	0.417	0.531	0.846	0.707	0.882	0.201	0.45	0.81
DAP067(2)	0.602	0.009	0.501	0.633	0.552	0.92	0.961	0.787	0.041	0.397	0.396	0.657	0.747	0.035	0.868	0.603	0.721	0.285	0.369	0.304	0.079	0.671	0.52	0.582	0.626	0.082	0.677	0.414	0.248	0.15
DAP121(1)	0.122	0.106	0.031	0.119	0.351	0.838	0.776	0.669	0.796	0.175	0.596	0.161	0.446	0.151	0.017	0.014	0.867	0.818	0.073	0.386	0.161	0.018	0.011	0.149	0.521	0.286	0.403	0.459	0.687	0.743
DAP114(1)	0.263	0.18	0.325	0.305	0.246	0.187	0.028	0.836	0.538	0.019	0	0.12	0.278	0.665	1	0.697	0.555	0.176	0.825	0.145	0.951	0.736	0.223	0.9	0.66	0.511	0.884	0.278	0.951	0.709
DAP043(2)	0.459	0.347	0.308	0.437	0.184	0.592	0.768	0.23	0.528	0.939	0.548	0.288	0.947	0.816	0.115	0.318	0.813	0.779	0.053	0.048	0.101	0.41	0.539	0.987	0.801	0.687	0.517	0.469	0.922	0.251
DAP018(1)	0.469	0.516	0.909	0.921	0.924	0.153	0.277	0.858	0.471	0.576	0.309	0.503	0.163	0.831	0.047	0.45	0.692	0.364	0.534	0.307	0.394	0.305	0.22	0.429	0.312	0.025	0.558	0.446	0.649	0.886
DAP094(1)	0.189	0.014	0.858	0.713	0.401	0.977	0.072	0.437	0.994	0.103	0.604	0.528	0.385	0.556	0.438	0.842	0.557	0.534	0.142	0.61	0.508	0.354	0.842	0.032	0.957	0.3	0.839	0.626	0.995	0.509
DAP042(2)	0.666	0.504	0.161	0.82	0.675	0.797	0.34	0.06	0.039	0.743	0.012	0.16	0.768	0.595	0.534	0.789	0.76	0.574	0.84	0.935	0.047	0.409	0.243	0.47	0.357	0.429	0.97	0.013	0.908	0.971
DAP065(1)	0.063	0.171	0.47	0.869	0.601	0.314	0.164	0.122	0.059	0.048	0.044	0.093	0.06	0.136	0.293	0.706	0.753	0.673	0.713	0.941	0.24	0.809	0.86	0.515	0.739	0.313	0.111	0.028	0.843	0
DAP028(2)	0.623	0.208	0.657	0.992	0.895	0.947	0.798	0.484	0.953	0.472	0.89	0.087	0.076	0.804	0.022	0.155	0.122	0.991	0.492	0.64	0.582	0.625	0.297	0.49	0.21	0.965	0.925	0.018	0.094	1
DAP092(1)	0.244	0.792	0.333	0.209	0.662	0.418	0.299	0.009	0.866	0.573	0.313	0.335	0.141	0.599	0.913	0.189	0.637	0.5	0.247	0.315	0.102	0.63	0.074	0.513	0.416	0.369	0.552	0.363	0.274	0.184
DAP084(2)	0.814	0.871	0.223	0.037	0.605	0.946	0.107	0.349	0.33	0.362	0.639	0.205	0.222	0.267	0.112	0.848	0.688	0.485	0.13	0.108	0.295	0.335	0.231	0.413	0.664	0.585	0.754	0.387	0.304	0.107
DAP105(1)	0.889	0.01	0.05	0.453	0.25	0.023	0.921	0.799	0.291	0.282	0.555	0.292	0.744	0.387	0.755	0.531	0.424	0.125	0.297	0.83	0.022	0.816	0.652	0.188	0.177	0.145	0.96	0.178	0.573	0.574
DAP011(3)	0.512	0.218	0.077	0.313	0.402	0.069	0.327	0.714	0.399	0.177	0.006	0.262	0.306	0.891	0.448	0.551	0.742	0.873	0.006	0.261	0.163	0.689	0.327	0.702	0.547	0.433	0.5	0.582	0.915	0.572
DAP051(1)	0.799	0.01	0.773	0.314	0.13	0.664	0.279	0.049	0.14	0.326	0.516	0.891	0.032	0.712	0.632	0.764	0.805	0.419	0.425	0.978	0.843	0.484	0.045	0.088	0.448	0.512	0.47	0.092	0.737	0.986
DAP110(1)	0.224	0.132	0.197	0.559	0.34	0.918	0.956	0.208	0.652	0.481	0.438	0.161	0.142	0.905	0	0.696	0.362	0.599	0.814	0.279	0.032	0.223	0.025	0.772	0.854	0.79	0.123	0.447	0.995	0.696
DAP077(1)	0.245	0.872	0.017	0.085	0.197	0.186	0.844	0.006	0.214	0.634	0.552	0.16	0.652	0.32	0.295	0.936	0.453	0.447	0.759	0.473	0.911	0.977	0.863	0.142	0.822	0.529	0.459	0.281	0.859	0.905
DAP053(1)	0.49	0.842	0.35	0.58	0.43	0.17	0.011	0.744	0.001	0.437	0.474	0.477	0.426	0.343	0.225	0.875	0.635	0.071	0.238	0.274	0.442	0.601	0.559	0.338	0.982	0.766	0.358	0.508	0.748	0.657
DAP058(1)	0.35	0.011	0.726	0.559	0.787	0.348	0.85	0.395	0.258	0.272	0.19	0.324	0.506	0.671	0.861	0.723	0.998	0.137	0.823	0.679	0.59	0.341	0.336	0.597	0.255	0.707	0.996	0.657	0.832	0.522
DAP106(1)	0.404	0.28	0.833	0.364	0.811	0.922	0.694	0.289	0.169	0.781	0.485	0.431	0.998	0.32	0.339	0.514	0.738	0.784	0.034	0.198	0.088	0.611	0.101	0.935	0.82	0.822	0.682	0.707	0.77	0.848
DAP056(3)	0.447	0.528	0.992	0.868	0.073	0.837	0.89	0.319	0.519	0.838	0.632	0.378	0.845	0.405	0.075	0.255	0.423	0.975	0.618	0.945	0.108	0.213	0.829	0.9	0.842	0.752	0.567	0.835	0.94	0.559
DAP030(1)	0.298	0.178	0.521	0.807	0.948	0.642	0.781	0.136	0.087	0.644	0.619	0.406	0.32	0.016	0.026	0.033	0.035	0.565	0.08	0.443	0.141	0.841	0.538	0.468	0.122	0.527	0.175	0.567	0.739	0.316
DAP093(1)	0.18	0.123	0.649	0.275	0.043	0.554	0.232	0.121	0.317	0.377	0.174	0.583	0.025	0.614	0.998	0.806	0.045	0.126	0.308	0.269	0.203	0.928	0.32	0.874	0.859	0.246	0.451	0.512	0.412	0.967
DAP059(2)	0.137	0.672	0.698	0.469	0.02	0.928	0.568	0.689	0.703	0.492	0.38	0.586	0.341	0.153	0.966	0.114	0.277	0.342	0.114	0.921	0.703	0.927	0.149	0.132	0.053	0.044	0.417	0.664	0.478	0.625
DAP086(2)	0.537	0.847	0	0.45	0.093	0.416	0.34	0.43	0.225	0.344	0.419	0.099	0.745	0.25	0.829	0.345	0.347	0.376	0.65	0.469	0.75	0.587	0.146	0.111	0.737	0.51	0.581	0.397	0.168	0.946
DAP104(1)	0.35	0.218	0.653	0.263	0.433	0.233	0.007	0.285	0.838	0.531	0.651	0.831	0.767	0.564	0.03	0.425	0.695	0.273	0.567	0.013	0.564	0.244	0.974	0.887	0.937	0.362	0.471	0.801	0.947	0.931
DAP082(1)	0.67	0.835	0.824	0.559	0.728	0.499	0.848	0.733	0.419	0.681	0.897	0.89	0.873	0.92	0.733	0.831	0.671	0.999	0.715	0.525	0.707	0.034	0.034	0.048	0.024	0.471	0.307	0.139	0.713	0.85

Trial	31	32	33	34	35	36	37	38	39	40	41	42	43	44	45	46	47	48	49	50	51	52	53	54	55	56	57	58	59	60
DAP022(2)	0.301	0.197	0.333	0.427	0.026	0.879	0.705	0.447	0.281	0.881	0.689	0.88	0.552	0.303	0.775	0.661	0.721	0.701	0.409	0.716	0.281	0.112	0.586	0.057	0.413	0.425	0.466	0.216	0.288	0.395
DAP080(1)	0.806	0.95	0.765	0.175	0.422	0.426	0.084	0.904	0.928	0.818	0.603	0.887	0.646	0.004	0.975	0.208	0.65	0.651	0.339	0.237	0.907	0.023	0.116	0.312	0.643	0.243	0.527	0.546	0.007	0.729
DAP075(2)	0.058	0.231	0.357	0.212	0.913	0.947	0.621	0.654	0.253	0.989	0.024	0.95	0.934	0.692	0.817	0.006	0.83	0.754	0.917	0.403	0.464	0.622	0.715	0.212	0.846	0.006	0.616	0.493	0.496	0.359
DAP087(1)	0.537	0.931	0.076	0.54	0.514	0.632	0.076	0.347	0.378	0.201	0.746	0.316	0.75	0.202	0.542	0.284	0.006	0.293	0.284	0.881	0.017	0.724	0.831	0.146	0.402	0.153	0.741	0.962	0.878	0.616
DAP101(1)	0.245	0.182	0.115	0.119	0.177	0.62	0.512	0.45	0.024	0.369	0.322	0.45	0.293	0.607	0.124	0.591	0.34	0.639	0.825	0.99	0.951	0.489	0.864	0.253	0.362	0.099	0.616	0.051	0.258	0.275
DAP125(1)	0.114	0.969	0.169	0.928	0.577	0.547	0.166	0.311	0.725	0.559	0.95	0.655	0.061	0.341	0.947	0.558	0.99	0.202	0.431	0.424	0.815	0.053	0.419	0.626	0.678	0.127	0.25	0.676	0.949	0.445
DAP016(1)	0.643	0.611	0.888	0.821	0.832	0.465	0.712	0.349	0.469	0.867	0.247	0.963	0.912	0.005	0.892	1	0.69	0.371	0.068	0.166	0.222	0.586	0.728	0.935	0.824	0.618	0.601	0.733	0.677	0.674
DAP072(1)	0.485	0.39	0.603	0.106	0.518	0.13	0.613	0.945	0.663	0.38	0.479	0.838	0.583	0.669	0.859	0.895	0.369	0.558	0.055	0.498	0.838	0.584	0.018	0.497	0.204	0.868	0.84	0.688	0.903	0.244
DAP060(1)	0.541	0.849	0.116	0.434	0.977	0.293	0.112	0.137	0.686	0.906	0.828	0.693	0.718	0.305	0.567	0.93	0.53	0.47	0.001	0.109	0.003	0.325	0.802	0.334	0.473	0.499	0.115	0.627	0.42	0.528
DAP088(1)	0.787	0.945	0.965	0.468	0.705	0.395	0.729	0.579	0.091	0.961	0.223	0.756	0.731	0.832	0.805	0.716	0.981	0.81	0.125	0.829	0.582	0.954	0.781	0.81	0.143	0.342	0.389	0.03	0.398	0.387
DAP102(1)	0.684	0.476	0.394	0.107	0.944	0.507	0.684	0.852	0.912	0.455	0.556	0.605	0.552	0.412	0.713	0.805	0.37	0.496	0.6	0.753	0.48	0.329	0.863	0.712	0.348	0.409	0.006	0.722	0.385	0.932
DAP001(3)	0.616	0.171	0.225	0.869	0.441	0.228	0.517	0.889	0.681	0.104	0.787	0.435	0.158	0.981	0.6	0.919	0.963	0.365	0.983	0.244	0.071	0.17	0.001	0.67	0.423	0.816	0.525	0.758	0.497	0.418
DAP126(1)	0.902	0.489	0.038	0.153	0.744	0.713	0.349	0.934	0.103	0.076	0.212	0.815	0.008	0.157	0.774	0.78	0.293	0.156	0.762	0.795	0.4	0.313	0.663	0.882	0.473	0.49	0.182	0.354	0.07	0.234
DAP108(1)	0.845	0.571	0.978	0.658	0.105	0.58	0.443	0.056	0.885	0.256	0.231	0.796	0.484	0.708	0.892	0.897	0.903	0.291	0.796	0.647	0.688	0.552	0.614	0.96	0.194	0.417	0.268	0.005	0.741	0.687
DAP095(2)	0.829	0.365	0.837	0.788	0.389	0.877	0.629	0.505	0.033	0.41	0.014	0.254	0.084	0.88	0.45	0.626	0.385	0.987	0.984	0.997	0.192	0.052	0.118	0.607	0.844	0.147	0.859	0.728	0.996	0.913
DAP084(1)	0.363	0.362	0.708	0.102	0.285	0.129	0.726	0.22	0.578	0.447	0.372	0.554	0.15	0.007	0.886	0.339	0.211	0.959	0.368	0.216	0.62	0.79	0.049	0.36	0.739	0.307	0.715	0.699	0.868	0.428
DAP081(1)	0.56	0.106	0.443	0.126	0.051	0.889	0.842	0.161	0.5	0.014	0.391	0.779	0.128	0.173	0.091	0.411	0.869	0.29	0.105	0.366	0.21	0.031	0.192	0.521	0.279	0.254	0.97	0.013	0.964	0.784
DAP075(1)	0.641	0.388	0.167	0.868	0.512	0.058	0.024	0.828	0.891	0.809	0.092	0.29	0.826	0.851	0.254	0.888	0.548	0.478	0.621	0.93	0.288	0.473	0.463	0.197	0.919	0.59	0.57	0.407	0.023	0.333
DAP031(2)	0.589	0.326	0.902	0.232	0.034	0.368	0.241	0.91	0.163	0.477	0.069	0.084	0.354	0.79	0.181	0.885	0.344	0.475	0.628	0.76	0.917	0.81	0.94	0.192	0.259	0.779	0.817	0.217	0.34	0.413
DAP054(1)	0.981	0.955	0.371	0.972	0.934	0.139	0.584	0.781	0.202	0.637	0.343	0.387	0.958	0.788	0.046	0.577	0.006	0.875	0.907	0.583	0.222	0.757	0.608	0.889	0.91	0.848	0.99	0.31	0.063	0.215
DAP069(2)	0.602	0.439	0.986	0.62	0.096	0.631	0.032	0.062	0.938	0.398	0.704	0.951	0.764	0.243	0.094	0.181	0.467	0.037	0.678	0.108	0.188	0.43	0.755	0.879	0.268	0.218	0.868	0.265	0.437	0.443
DAP091(1)	0.408	0.792	0.993	0.411	0.498	0.049	0.396	0.254	0.701	0.721	0.961	0.021	0.366	0.259	0.564	0.351	0.14	0.93	0.378	0.861	0.708	0.546	0.858	0.117	0.245	0.884	0.142	0.312	0.744	0.796
DAP096(2)	0.9	0.308	0.602	0.294	0.065	0.53	0.777	0.686	0.505	0.797	0.197	0.788	0.412	0.3	0.261	0.225	0.708	0.196	0.87	0.692	0.016	0.929	0.714	0.231	0.642	0.153	0.324	0.876	0.989	0.41
DAP090(1)	0.954	0.009	0.205	0.518	0.039	0.084	0.648	0.051	0.47	0.71	0.801	0.776	0.794	0.918	0.167	0.232	0.639	0.745	0.387	0.567	0.453	0.364	0.464	0.2	0.375	0.987	0.144	0.378	0.186	0.512
DAP097(2)	0.756	0.766	0.839	0.841	0.701	0.235	0.135	0.051	0.598	0.556	0.946	0.962	0.474	0.078	0.854	0.349	0.059	0.849	0.615	0.066	0.22	0.768	0.064	0.212	0.03	0.143	0.446	0.117	0.55	0.454
DAP048(2)	0.651	0.436	0.499	0.114	0.549	0.162	0.248	0.238	0.3	0.72	0.194	0.445	0.248	0.867	0.358	0.799	0.235	0.645	0.632	0.801	0.507	0.009	0.243	0.073	0.617	0.267	0.651	0.124	0.788	0.106
DAP109(1)	0.653	0.641	0.387	0.156	0.03	0.079	0.455	0.589	0.81	0.326	0.689	0.123	0.623	0.468	0.114	0.308	0.574	0.304	0.737	0.533	0.816	0.911	0.281	0.841	0.112	0.271	0.567	0.698	0.875	0.302
DAP050(1)	0.707	0.324	0.925	0.914	0.284	0.227	0.855	0.596	0.617	0.603	0.828	0.182	0.916	0.023	0.815	0.54	0.206	0.616	0.84	0.849	0.756	0.669	0.831	0.923	0.884	0.142	0.937	0.186	0.499	0.418
DAP057(3)	0.803	0.012	0.768	0.868	0.675	0.514	0.549	0.896	0.869	0.168	0.239	0.869	0.789	0.173	0.049	0.066	0.77	0.71	0.409	0.025	0.749	0.269	0.444	0.47	0.613	0.688	0.755	0.781	0.246	0.436
DAP095(1)	0.843	0.743	0.855	0.792	0.447	0.555	0.682	0.27	0.152	0.767	0.501	0.82	0.909	0.858	0.267	0.591	0.063	0.325	0.668	0.728	0.958	0.805	0.424	0.03	0.521	0.56	0.116	0.842	0.97	0.599
DAP067(1)	0.974	0.287	0.971	0.405	0.036	0.292	0.833	0.932	0.7	0.061	0.323	0.305	0.403	0.329	0.693	0.797	0.514	0.689	0.308	0.062	0.432	0.559	0.866	0.69	0.038	0.878	0.581	0.279	0.774	0.322
DAP014(2)	0.703	0.808	0.822	0.374	0.69	0.165	0.507	0.984	0.776	0.981	0.784	0.939	0.123	0.215	0.083	0.309	0.878	0.105	0.263	0.831	0.762	0.364	0.278	0.067	0.27	0.528	0.58	0.536	0.486	0.308
DAP123(1)	0.967	0.138	0.785	0.159	0.273	0.97	0.834	0.681	0.618	0.547	0.777	0.75	0.501	0.48	0.149	0.62	0.265	0.17	0.474	0.221	0.312	0.863	0.896	0.224	0.407	0.045	0.219	0.44	0.112	0.446
DAP067(2)	0.649	0.414	0.246	0.103	0.346	0.289	0.455	0.903	0.784	0.312	0.126	0.876	0.179	0.443	0.755	0.63	0.613	0.525	0.867	0.006	0.279	0.498	0.747	0.919	0.936	0.929	0.681	0.25	0.41	0.046
DAP121(1)	0.174	0.314	0.097	0.309	0.939	0.928	0.007	0.4	0.575	0.027	0.058	0.672	0.291	0.156	0.047	0.652	0.341	0.137	0.313	0.425	0.554	0.426	0.6	0.854	0.118	0.298	0.956	0.523	0.478	0.645
DAP114(1)	0.052	0.017	0.165	0.529	0.417	0.958	0.606	0.463	0.494	0.84	0.094	0.051	0.654	0.149	0.313	0.128	0.055	0.539	0.765	0.009	0.938	0.89	0.298	0.458	0.002	0.439	0.881	0.36	0.978	0.236
DAP043(2)	0.807	0.936	0.89	0.732	0.703	0.79	0.994	0.57	0.19	0.608	0.2	0.43	0.289	0.198	0.75	0.751	0.952	0.344	0.53	0.839	0.436	0.421	0.627	0.428	0.201	0.32	0.989	0.895	0.882	0.708
DAP018(1)	0.973	0.707	0.401	0.344	0.903	0.412	0.464	0.541	0.385	0.351	0.181	0.714	0.775	0.092	0.264	0.9	0.314	0.322	0.697	0.529	0.736	0.737	0.95	0.497	0.48	0.402	0.142	0.754	0.331	0.737
DAP094(1)	0.404	0.933	0.279	0.378	0.24	0.992	0.641	0.485	0.004	0.179	0.845	0.923	0.221	0.069	0.076	0.345	0.018	0.507	0.253	0.379	0.717	0.436	0.154	0.096	0.148	0.157	0.807	0.594	0.996	0.558
DAP042(2)	0.535	0.868	0.031	0.701	0.09	0.935	0.095	0.543	0.948	0.902	0.402	0.123	0.809	0.557	0.496	0.555	0.096	0.249	0.742	0.817	0.456	0.324	0.815	0.672	0.618	0.624	0.118	0.12	0.421	0.258
DAP065(1)	0.869	0.876	0.665	0.241	0.28	0.896	0.39	0.538	0.171	0.109	0.096	0.333	0.513	0.388	0.686	0.789	0.816	0.962	0.772	0.896	0.71	0.959	0.957	0.961	0.768	0.648	0.729	0.569	0.65	0.911
DAP028(2)	0.081	0.031	0.577	0.974	0.783	0.657	0.902	0.799	0.554	0.568	0.491	0.473	0.971	0.781	0.627	0.913	0.749	0.045	0.997	0.509	0.662	0.934	0.52	0.631	0.799	0.222	0.727	0.544	0.812	0.62
DAP092(1)	0.025	0.165	0.839	0.798	0.642	0.981	0.536	0.915	0.938	0.879	0.494	0.871	0.411	0.848	0.351	0.544	0.614	0.9	0.217	0.529	0.484	0.56	0.783	0.897	0.648	0.58	0.627	0.937	0.959	0.448
DAP084(2)	0.735	0.808	0.338	0.941	0.855	0.599	0.885	0.585	0.77	0.658	0.798	0.307	0.095	0.41	0.749	0.249	0.645	0.926	0.149	0.108	0.435	0.997	0.753	0.537	0.298	0.062	0.646	0.156	0.046	0.127
DAP105(1)	0.054	0.772	0.39	0.377	0.363	0.281	0.723	0.31	0.54	0.866	0.671	0.743	0.024	0.774	0.287	0.987	0.663	0.344	0.541	0.296	0.733	0.63	0.192	0.188	0.533	0.016	0.963	0.074	0.215	0.15
DAP011(3)	0.695	0.171	0.877	0.243	0.424	0.561	0.627	0.85	0.132	0.981	0.165	0.392	0.534	0.987	0.408	0.215	0.1	0.432	0.252	0.485	0.329	0.972	0.417	0.07	0.236	0.407	0.823	0.757	0.35	0.647
DAP051(1)	0.179	0.696	0.783	0.65	0.804	0.93	0.064	0.776	0.381	0.252	0.8	0.558	0.975	0.003	0.898	0.235	0.789	0.138	0.856	0.095	0.745	0.178	0.058	0.571	0.624	0.369	0.454	0.576	0.538	0.202
DAP110(1)	0.778	0.949	0.45	0.027	0.944	0.245	0.108	0.798	0.236	0	0.364	0.239	0.069	0.891	0.713	0.892	0.974	0.662	0.686	0.64	0.049	0.549	0.058	0.776	0.816	0.107	0.196	0.377	0.274	0.006
DAP077(1)	0.745	0.78	0.48	0.776	0.242	0.542	0.892	0.828	0.382	0.222	0.639	0.755	0.437	0.301	0.305	0.87	0.357	0.116	0.138	0.368	0.761	0.339	0.724	0.102	0.501	0.901	0.426	0.432	0.534	0.433
DAP053(1)	0.717	0.248	0.599	0.183	0.681	0.604	0.622	0.959	0.575	0.377	0.731	0.512	0.542	0.024	0.525	0.749	0.026	0.021	0.756	0.311	0.583	0.292	0.596	0.165	0.924	0.057	0.649	0.879	0.042	0.944
DAP058(1)	0.751	0.276	0.987	0.988	0.051	0.808	0.946	0.578	0.894	0.144	0.413	0.449	0.228	0.131	0.766	0.442	0.883	0.6	0.697	0.574	0.166	0.602	0.723	0.123	0.468	0.731	0.642	0.169	0.516	0.211
DAP106(1)	0.46	0.351	0.872	0.952	0.394	0.825	0.848	0.315	0.301	0.698	0.841	0.485	0.508	0.637	0.178	0.457	0.429	0.634	0.281	0.337	0.954	0.597	0.863	0.202	0.406	0.742	0.788	0.965	0.039	0.44
DAP056(3)	0.591	0.631	0.996	0.768	0.735	0.931	0.554	0.265	0.229	0.984	0.717	0.692	0.409	0.664	0.804	0.183	0.797	0.827	0.652	0.327	0.228	0.369	0.573	0.554	0.288	0.348	0.751	0.561	0.195	0.53
DAP030(1)	0.392	0.686	0.258	0.076	0.359	0.435	0.405	0.435	0.798	0.927	0.729	0.753	0.91	0.852	0.706	0.31	0.122	0.272	0.882	0.272	0.588	0.85	0.531	0.758	0.321	0.944	0.952	0.935	0.55	0.253
DAP093(1)	0.509	0.308	0.541	0.688	0.303	0.135	0.822	0.719	0.285	0.217	0.803	0.262	0.701	0.987	0.387	0.045	0.827	0.877	0.716	0.214	0.276	0.769	0.911	0.68	0.058	0.787	0.811	0.527	0.625	0.273
DAP059(2)	0.36	0.001	0.09	0.555	0.255	0.64	1	0.357	0.467	0.396	0.027	0.171	0.153	0.78	0.369	0.703	0.965	0.919	0.505	0.374	0.632	0.094	0.344	0.99	0.904	0.056	0.73	0.708	0.388	0.898
DAP086(2)	0.628	0.617	0.748	0.687	0.45	0.213	0.795	0.012	0.95	0.486	0.71	0.755	0.578	0.462	0.282	0.57	0.871	0.686	0.205	0.7	0.282	0.437	0.883	0.416	0.06	0.722	0.194	0.633	0.46	0.873
DAP104(1)	0.138	0.354	0.964	0.393	0.674	0.364	0.229	0.823	0.761	0.949	0.076	0.59	0.118	0.091	0.679	0.279	0.354	0.961	0.296	0.223	0.48	0.82	0.173	0.197	0.212	0.487	0.684	0.587	0.033	0.853
DAP082(1)	0.466	0.459	0.15	0.948	0.89	0.581	0.239	0.191	0.512	0.341	0.528	0.356	0.414	0.317	0.629	0.707	0.509	0.353	0.378	0.409	0.631	0.422	0.792	0.405	0.531	0.899	0.736	0.655	0.505	0.87

The following is a full table of p-values obtained from running the modified experiment to test for linear dependencies using PCMCI. This modified experiment ignores missing data (which restricts the maximum lag length tested for each trial depending on sample size) and isolates the effects of increases in $PM_{2.5}$. Each row shows p-values (rounded to 3.d.p.) for a given trial obtained from tests at all lag lengths up to the maximum possible lag length. P-values are also colour-coded as follows:

- $p < 0.05$: Green.
- $0.05 \leq p < 0.1$: Yellow.
- $0.1 \leq p$: Red.

Trial	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31	32	33	34	35	36	37	38	39	40	41	42	43	44	45	46	47	48
DAP022(2)	0.688	0.497	0.081	0.741	0.858	0.906	0.607	0.649	0.066	0.924	0.867	0.621	0.371	0.014	0.713	0.001	0.054	0.201	0.777	0.456																												
DAP080(1)	0.199	0.388	0.721	0.551	0.039	0.062	0.048	0.013	0.008	0.01	0.27	0.241	0.459	0.294	0.506	0.443	0.14	0.483	0.189																													
DAP075(2)	0.055	0.297	0.159	0.357	0.125	0.633	0.904	0.565	0.917	0.945	0.367	0.632	0.724	0.849	0.194	0.868	0.365	0.092	0.713	0.509	0.725	0.801	0.632	0.072	0.386	0.401	0.153	0.822	0.391	0.036	0.549	0.176	0.321	0.129	0.671											
DAP087(1)	0.537	0.305	0.495	0.379	0.453	0.457	0.458	0.478	0.46	0.448	0.505	0.413	0.44	0.381	0.098	0.496	0.277	0.914	0.3	0.9	0.262	0.798																								
DAP101(1)	0.516	0.879	0.364	0.407	0.589	0.224	0.063	0.073	0.687	0.008	0.132	0.321	0.623	0.983	0.358	0.177	0.976	0.866	0.675																													
DAP125(1)	0.413	0.137	0.672	0.887	0.49	0.025	0.456	0.773	0.011	0.729	0.264	0.85	0.381	0.063	0.516																															
DAP016(1)	0.736	0.481	0.952	0.206	0.125	0.371	0.453	0.465	0.599	0.9	0.863	0.643	0.544	0.291	0.351	0.614	0.825	0.873	0.264	0.01	0.252	0.694	0.615	0.113	0.526	0.729	0.893	0.89	0.163	0.478														
DAP072(1)	0.457	0.477	0.325	0.544	0.857	0.501	0.055	0.284	0.827	0.888	0.584	0.117	0.506	0.034	0.554																															
DAP060(1)	0.613	0.757	0.474	0.324	0.701	0.105	0.509	0.27	0.563	0.416	0.171	0.353	0.723	0.758	0.474	0.686	0.063	0.271	0.499	0.211	0.482	0.226	0.418	0.705																				
DAP088(1)	0.766	0.18	0.721	0.805	0.866	0.659	0.656																																							
DAP102(1)	0.287	0.579	0.565	0.195	0.022																																								
DAP001(3)	0.603	0.137	0.591	0.723	0.039	0.186	0.334	0.068	0.254	0.094	0.361	0.446	0.248	0.184	0.362	0.29	0.236																													
DAP126(1)	0.447	0.16	0.44	0.211	0.108	0.707	0.253	0.312	0.222	0.738	0.035	0.574	0.491	0.408	0.592	0.846	0.453	0.225	0.445	0.774	0.209	0.273	0.571	0.894	0.658	0.34	0.752	0.49	0.043															
DAP108(1)	0.086	0.895	0.374	0.017	0.016	0.742	0.663	0.468	0.462	0.279	0.477	0.845	0.188	0.468	0.902	0.852	0.592	0.108																													
DAP095(2)	0.642	0.528	0.311	0.899	0.813	0.107	0.806	0.312	0.04	0.415	0.509	0.289	0.338	0.982	0.895	0.224	0.039	0.549	0.82	0.193	0.086	0.38	0.527	0.472	0.969	0.741	0.038	0.543	0.353	0.314	0.394	0.608	0.382	0.296											
DAP084(1)	0.916	0.969	0.011	0.652	0.765	0.314	0.631	0.609	0.62	0.761	0.844	0.325	0.107	0.74	0.933	0.613	0.704	0.112																													
DAP081(1)	0.375	0.613	0.221	0.602	0.96	0.246	0.745	0.782	0.795	0.579	0.876	0.737	0.153	0.821																															
DAP075(1)	0.139	0.734	0.977	0.073	0.963	0.673	0.972	0.63	0.586	0.443	0.067	0.238	0.304	0.314	0.652	0.425	0.181	0.335	0.855	0.397	0.405	0.117	0.449																						
DAP031(2)	0.572	0.81	0.333	0.679	0.349	0.904	0.154	0.993	0.865	0.739	0.236	0.975	0.216	0.371	0.714	0.804	0.578	0.281	0.16	0.346	0.601	0.803	0.957																						
DAP054(1)	0.328	0.552	0.776	0.623	0.906	0.241	0.59	0.488	0.949	0.479	0.849	0.095	0.719	0.613	0.925	0.851	0.286	0.149																													
DAP069(2)	0.667	0.486	0.938	0.877	0.436	0.348	0.094	0.507	0.597	0.049	0.866	0.531	0.595																																	
DAP091(1)	0.031	0.24	0.733	0.03	0.248	0.377	0.086	0.822	0.129	0.712	0.001	0.658	0.191	0.177	0.49	0.249	0.293	0.175	0.855	0.813	0.079	0.539	0.117	0.976	0.079	0.201	0.386																		
DAP096(2)	0.856	0.963	0.53	0.691	0.954	0.906	0.378	0.84	0.975	0.064	0.158	0.312	0.445	0.818	0.04	0.712	0.11	0.7	0.019	0.466	0.74	0.542	0.491	0.117	0.266																				
DAP090(1)	0.856	0.436	0.002	0.301	0.961	0.489	0.491	0.956	0.118	0.789	0.214	0.303	0.443	0.719	0.914	0.218	0.652	0.849	0.358	0.355	0.973																								
DAP097(2)	0.695	0.043	0.723	0.017	0.038	0.218	0.629	0.959	0.228	0.98	0.303	0.022	0.224	0.62	0.25	0.072	0.256	0.222	0.803	0.446	0.549																								
DAP048(2)	0.468	0.874	0.706	0.233	0.879	0.858	0.832	0.096	0.857	0.308	0.224	0.473	0.963	0.519	0.757	0.977	0.517	0.901	0.932	0.11	0.812	0.165	0.493	0.379	0.219	0.455	0.972																		
DAP109(1)	0.341	0.059	0.808	0.241	0.194	0.423	0.845	0.372	0.52	0.112	0.297	0.763	0.149	0.67	0.611	0.545	0.039	0.711	0.447	0.257	0.571	0.368	0.763	0.708	0.407	0.503	0.805																		
DAP050(1)	0.199	0.116	0.422	0.2	0.7	0.424	0.443	0.023	0.264	0.503	0.84	0.647	0.661	0.331	0.741	0.038	0.483	0.286	0.436	0.952	0.217	0.824	0.542	0.577	0.984	0.194	0.038	0.131	0.991	0.045	0.088	0.588	0.066	0.597	0.746	0.844	0.242							
DAP057(3)	0.284	0.714	0.042	0.036	0.575	0.537	0.513	0.969	0.149	0.539	0.469	0.959	0.111	0.638	0.42	0.879	0.476	0.455	0.477	0.322	0.182	0.707	0.381	0.519	0.716	0.518	0.815	0.165	0.116	0.276														
DAP095(1)	0.288	0.885	0.001	0.63	0.826	0.176	0.918	0.848	0.43	0.243	0.583	0.728	0.992	0.673	0.146	0.989	0.272	0.649	0.676	0.117	0.639	0.745	0.666	0.92	0.698	0.542																		
DAP067(1)	0.483	0.461	0.326	0.276	0.114	0.105	0.02	0.235	0.615	0.087	0.154	0.292	0.015	0.535	0.136	0.159	0.887																												
DAP014(2)	0.364	0.38	0.089	0.109	0.893	0.577	0.394	0.084	0.049	0.041	0.342	0.612	0.528	0.98	0.331	0.817	0.043																												
DAP123(1)	0.546	0.173	0.747	0.28	0.353	0.135	0.183	0.139	0.678	0.146	0.779	0.659	0.292	0.877	0.863	0.413	0.335	0.96	0.704	0.154	0.731	0.609	0.517																						
DAP067(2)	0.854	0.994	0.284	0.139	0.038	0.693	0.562	0.85	0.283	0.861	0.26	0.26	0.575	0.865	0.443	0.252	0.413	0.389	0.561	0.779	0.992	0.612	0.453	0.163	0.128	0.275																		
DAP121(1)	0.489	0.308	0.936	0.3	0.927	0.2	0.715	0.461	0.133	0.248	0.846	0.517	0.706	0.191	0.434	0.575	0.994	0.23	0.034	0.531	0.945	0.491	0.336	0.359	0.965	0.763	0.512	0.864																
DAP114(1)	0.218	0.476	0.874	0.857	0.214	0.709	0.368	0.167	0.536	0.902	0.961	0.521	0.627	0.473	0.989	0.184	0.682	0.071																												
DAP043(2)	0.147	0.691	0.352	0.111	0.186	0.01	0.03	0.417	0.379	0.169	0.359	0.374	0.91	0.002	0.318																													
DAP018(1)	0.77	0.435	0.405	0.491	0.086	0.693	0.897	0.921	0.504	0.74	0.237	0.374	0.181	0.551	0.857	0.056	0.035	0.642																												
DAP094(1)	0.825	0.973	0.931	0.91	0.416	0.731	0.416	0.407	0.194	0.816	0.724	0.728	0.319	0.815	0.786	0.589	0.432	0.782	0.898	0.242	0.642	0.824	0.568	0.714	0.674	0.728	0.496	0.424	0.555	0.031	0.607														
DAP042(2)	0.678	0.279	0.978	0.441	0.699	0.424	0.654	0.802	0.919	0.799	0.913	0.899	0.302	0.34	0.239	0.299	0.497	0.908	0.4	0.642	0.857																								
DAP065(1)	0.639	0.176	0.688	0.904	0.461	0.076	0.175	0.08	0.275	0.984	0.16	0.088	0.948	0.389	0.923																													
DAP028(2)	0.611	0.891	0.664	0.36	0.099	0.333	0.165	0.184	0.516	0.084	0.079	0.777	0.655	0.451	0.204	0.078	0.943	0.155	0.363	0.999	0.149	0.86	0.563	0.802	0.566	0.174	0.224	0.462	0.021	0.514	0.437	0.923	0.57	0.277	0.058	0.837	0.472	0.493	0.048	0.796	0.671	0.305	0.211	0.087	0.096	0.897	0.811	0.654
DAP092(1)	0.241	0.172	0.209	0.434	0.314	0.969	0.455	0.383	0.44	0.331	0.155	0.186	0.211	0.313	0.364																															
DAP084(2)	0.024	0.247	0.988	0.891	0.523	0.237	0.644	0.42	0.636	0.874	0.102	0.24	0.953	0.102	0.397	0.894	0.865	0.266	0.214	0.75	0.994	0.205	0.026	0.691																				
DAP105(1)	0.304	0.597	0.309	0.121	0.763	0.853	0.193	0.531	0.733	0.199	0.674	0.462	0.699	0.34	0.081	0.211	0.627	0.365	0.789	0.337	0.441	0.307	0.151	0.8	0.434																				
DAP011(3)	0.806	0.479	0.747	0.65	0.284	0.717	0.385	0.543	0.216	0.288	0.064	0.918	0.704	0.449	0.869	0.331	0.223	0.337	0.791	0.312																										
DAP051(1)	0.584	0.719	0.706	0.7	0.999	0.815	0.759	0.56	0.654	0.51	0.13	0.791	0.899	0.586	0.87	0.585	0.707	0.873	0.979	0.165	0.607	0.525	0.957	0.829	0.861	0.62	0.139	0.67	0.795	0.996	0.464	0.993												
DAP110(1)	0.999	0.129	0.654	0.618	0.373	0.111	0.791	0.239	0.521	0.678	0.959	0.599	0.951	0.826	0.506	0.16	0.504	0.457	0.023	0.42																										
DAP077(1)	0.711	0.473	0.551	0.206	0.965	0.788	0.982	0.825	0.834	0.946	0.948	0.635	0.045	0.177	0.523	0.786	0.748	0.75	0.948	0.863																										
DAP053(1)	0.719	0.176	0.259	0.485	0.775	0.337	0.586	0.359	0.777	0.93	0.051	0.039	0.057	0.773	0.13	0.404	0.628	0.009	0.429	0.633	0.372	0.741																						
DAP058(1)	0.819	0.28	0.64	0.51	0.014	0.67	0.426	0.22	0.341	0.294	0.899	0.922	0.101	0.328	0.298	0.335	0.533	0.144	0.515	0.744	0.158	0.906	0.921																				
DAP106(1)	0.335	0.294	0.049	0.702	0.839	0.148	0.515	0.788	0.502	0.124	0.525	0.752	0.663	0.478	0.868	0.669	0.382	0.489	0.203	0.913	0.85	0.956																				
DAP056(3)	0.566	0.536	0.361	0.272	0.332	0.881	0.522	0.658	0.824	0.878	0.676	0.258	0.258	0.363	0.478	0.796	0.726	0.81	0.385	0.153	0.674	0.919	0.449	0.083	0.496	0.554																
DAP030(1)	0.335	0.359	0.052	0.214	0.624	0.373	0.026	0.154	0.51	0.332	0.056	0.373	0.749	0.697	0.47	0.922	0.421	0.976	0.376	0.264	0.666	0.569	0.189	0.367	0.368	0.095	0.754	0.689														
DAP093(1)	0.123	0.948	0.31	0.446	0.856	0.677	0.398	0.186	0.872	0.806	0.736	0.625	0.978	0.988	0.899	0.468	0.325	0.467	0.791																										
DAP059(2)	0.424	0.366	0.096	0.021	0.091	0.474	0.259	0.92	0.078	0.258	0.236	0.059	0.029	0.176	0.218	0.073	0.008	0.069	0.56	0.378	0.014	0.746	0.769	0.328	0.064	0.002	0.088																
DAP086(2)	0.646	0.772	0.558	0.29	0.697	0.135	0.905	0.321	0.34	0.978	0.986	0.899	0.494	0.805	0.862	0.579	0.388	0.97	0.925	0.809	0.422	0.281	0.908	0.758	0.748	0.864	0.754	0.117														
DAP104(1)	0.907	0.512	0.828	0.818	0.163	0.531	0.244	0.891	0.023	0.609	0.11	0.115	0.873	0.884	0.701	0.751	0.351	0.03	0.955	0.936	0.355	0.919	0.526	0.459	0.587	0.979	0.661	0.848	0.27	0.518												
DAP082(1)	0.824	0.612	0.597	0.722	0.866	0.896	0.785	0.481	0.525	0.74	0.908	0.556	0.784	0.958	0.893	0.995	0.877	0.934	0.803	0.977	0.918	0.856	0.866	0.995	0.811	0.438																

Appendix D

Full Results of PCMCI for Non-Linear Dependencies

The following is a full table of p-values obtained from running PCMCI with a nearest-neighbour based estimator for conditional mutual information to identify non-linear dependencies. Each row shows p-values (rounded to 3.d.p.) for a given trial obtained from tests at lag lengths of 1, 5, 10, 15, 30, 45 and 60 minutes. P-values are also colour-coded as follows:

- $p < 0.05$: Green.
- $0.05 \leq p < 0.1$: Yellow.
- $0.1 \leq p$: Red.

Trial	1	5	10	15	30	45	60
DAP001(3)	0.008	0.608	0.669	0	0	0.403	0.738
DAP011(3)	0	0	0	0	0	0.02	0.029
DAP014(2)	0.295	0	0	0.029	0.076	0.007	0.289
DAP016(1)	0.025	0	0	0	0.143	0.064	0.02
DAP018(1)	0.777	0.073	0.294	0.996	0.837	0.572	0.27
DAP022(2)	0.059	0.003	0.139	0.471	0.514	0.097	0.37
DAP022(2)	0.095	0.003	0.139	0.471	0.514	0.097	0.37
DAP028(2)	0.035	0.093	0	0	0.055	0.343	0.415
DAP030(1)	0.724	0.035	0.002	0	0.091	0.162	0.131
DAP031(2)	0.943	1	0.293	0.992	0.03	0	0
DAP042(2)	0.015	0	0	0	0.001	0	0.162
DAP043(2)	0.097	0.001	0.061	0.177	0.675	0.961	0.703
DAP048(2)	0.654	0	0	0.037	0.028	0.098	0.061
DAP050(1)	0.004	0	0	0	0.188	0.377	0.769
DAP051(1)	0.254	0	0	0	0	0	0
DAP053(1)	0.968	0.948	0.697	0.985	0.963	0.442	0.321
DAP054(1)	1	0	0	0	0.003	0	0.571
DAP056(3)	1	0	0	0	0	0	0
DAP057(3)	0.03	0.02	0.016	0.147	0.321	0.629	0.798
DAP058(1)	0.516	0.004	0.012	0.001	0	0.493	0.072
DAP059(2)	0	0	0	0	0	0.011	0
DAP060(1)	0.055	0	0	0.001	0.753	0.594	0.855
DAP065(1)	0.071	0	0.01	0.02	0.027	0.418	0.351
DAP067(1)	0.209	0.291	0.197	0.65	0.8	0.36	0.729
DAP067(2)	0.063	0	0	0	0	0.015	0.021
DAP069(2)	0.15	0.113	0.603	0.285	0.073	0.251	0.382
DAP072(1)	0.256	0	0	0	0.036	0.005	0.01
DAP075(1)	0.933	0	0	0	0	0.027	0.774
DAP075(2)	1	0.175	0.045	0.002	0.496	0.824	0.654
DAP077(1)	0.053	0	0	0.013	0	0.002	0.457
DAP080(1)	0.34	0.001	0	0.013	0.888	0	0.197
DAP081(1)	0.018	0	0	0	0.018	0.009	0
DAP082(1)	0.108	0	0	0.015	0.262	0.093	0.054
DAP084(1)	0	0	0	0	0	0	0
DAP084(2)	0.025	0	0	0	0	0.045	0.022
DAP086(2)	0	0	0	0	0	0.16	0.359
DAP087(1)	0	0.039	0	0.024	0.006	0.355	0.223
DAP088(1)	0.087	0.837	0.537	0.158	0.185	0.776	0.727
DAP090(1)	0.54	0	0	0	0	0	0
DAP091(1)	0	0	0.017	0.102	0.727	0.887	0.742
DAP092(1)	0.094	0	0	0	0.222	0.079	0.021
DAP093(1)	0.024	0	0	0	0.001	0.031	0.474
DAP094(1)	0.395	0.002	0	0	0.207	0.063	0
DAP095(1)	0.011	0.003	0.003	0.137	0.204	0.231	0.41
DAP095(2)	0.02	0.863	0.845	0.306	0.804	0.305	0.13
DAP096(2)	0.035	0	0	0	0.018	0	0.001
DAP097(2)	0	0	0	0	0.024	0.018	0
DAP101(1)	0	0.003	0	0.006	0.015	0	0
DAP102(1)	0.38	0.912	0.72	0.935	0.691	0.932	0.744
DAP104(1)	0.612	0.015	0	0.002	0	0	0
DAP105(1)	0.014	0.042	0.142	0.244	0.537	0.497	0.321
DAP106(1)	0.114	0.795	0.069	1	0.568	0.191	0.644
DAP108(1)	1	0.004	0.006	0.182	0.002	0.594	0.647
DAP109(1)	0.824	0.002	0.413	0.072	0.077	0.048	0.593
DAP110(1)	0.1	0	0	0	0.057	0	0
DAP114(1)	0.002	0	0	0	0	0	0
DAP121(1)	0.141	0.001	0	0	0	0	0.926
DAP123(1)	0.233	0	0.003	0	0.003	0.002	0.612
DAP125(1)	0.116	0	0	0	0	0	0
DAP126(1)	0.997	0.551	0	0.025	0	0.573	0.648

Bibliography

[1] D.K. Arvind, C.A. Bates, D.J. Fischer, and J Mann. A sensor data collection environment for clinical trials investigating health effects of airborne pollution. *2018 IEEE EMBS International Conference on Biomedical & Health Informatics*, 2018.

[2] D.K. Arvind, C.A. Bates, D.J. Fischer, and J Mann. Spatially-resolved estimation of personal dosage of airborne particulates for ambulatory subjects using wearable sensors. *2018 IEEE EMBS International Conference on Biomedical & Health Informatics*, 2018.

[3] A. Bates, M. J. Ling, J. Mann, and D. K. Arvind. Respiratory rate and flow waveform estimation from tri-axial accelerometer data. *2010 International Conference on Body Sensor Networks*, pages 144–150, 2010.

[4] MA Bind. Causal modeling in environmental health. *Annu Rev Public Health*, 40:23–43, 2019.

[5] LA Cox. Do causal concentration-response functions exist? a critical review of associational and causal relations between fine particulate matter and mortality. *Crit. Rev. Toxicol.*, 47:603–31, 2017.

[6] LA Cox. Modernizing the bradford hill criteria for assessing causal relationships in observational data. *Crit. Rev. Toxicol*, 15:1–31, 2018.

[7] LA Cox. Improving causal determination. *Global Epidemiology*, 1, 2019.

[8] C Diks and V Panchenko. A note on the hiemstra-jones test for granger noncausality. *Studies in Nonlinear Dynamics and Econometrics*, 9, 2005.

[9] C Diks and V Panchenko. A new statistic and practical guidelines for nonparametric granger causality testing. *Journal of Economic Dynamics & Control*, 30:1647–1669, 2006.

[10] Alex Carll et al. Inhaled ambient-level traffic-derived particulates decrease cardiac vagal influence and baroreflexes and increase arrhythmia in a rat model of metabolic syndrome. *Particle and Fibre Toxicology*, 14, 2017.

[11] AM Padula et al. Exposure to traffic-related air pollution during pregnancy and term low birth weight: Estimation of causal associations in a semiparametric model. *Am J Epidemiol*, 176:815–824, 2012.

[12] CWS Chen et al. Causality test of ambient fine particles and human influenza in taiwan: Age group-specific disparity and geographic heterogeneity. *Environmental International*, 111:354–361, 2018.

[13] DW Dockery et al. An association between air pollution and mortality in six u.s. cities. *The New England Journal of Medicine*, 329, 1993.

[14] J. Mann et al. Simultaneous activity and respiratory monitoring using an accelerometer. *2011 International Conference on Body Sensor Networks*, pages 139–143, 2011.

[15] J Runge et al. Detecting and quantifying causal associations in large nonlinear time series datasets. *Sci. Adv.*, 5, 2019.

[16] J Schwartz et al. Particulate air pollution and hospital emergency room visits for asthma in seattle. *Am. Rev. Respir. Dis.*, 147:826–31, 1993.

[17] LA Cox et al. Temperature, not fine particulate matter (pm 2.5), is causally associated with short-term acute daily mortality rates: results from one hundred united states cities. *Dose Response*, 11:319–43, 2012.

[18] LA Cox et al. Causal versus spurious spatial exposure-response associations in health risk analysis. *Critical Reviews in Toxicology*, 43:26–38, 2013.

[19] Nilsa Loyo-Berrios et al. Air pollution sources and childhood asthma attacks in cataño, puerto rico. *American Journal of Epidemiology*, 165:927–35, 2007.

[20] Olivier Laurent et al. Air pollution, asthma attacks, and socioeconomic deprivation: A small-area case-crossover study. *American Journal of Epidemiology*, 168:58–65, 2008.

[21] P Basu et al. A nonparametric test for stationarity based on local fourier analysis. *IEEE International Conference on Acoustics, Speech and Signal Processing*, pages 3005–3008, 2009.

[22] S Kesten et al. Respiratory rate during acute asthma. *Chest*, 97:58–62, 1990.

[23] S Selvin et al. Ecologic regression analysis and the study of the influence of air quality on mortality. *Environmental Health Perspectives*, 54:333–340, 1984.

[24] Shaoping Yang et al. Ambient air pollution the risk of stillbirth: A prospective birth cohort study in wuhan, china. *International Journal of Hygiene and Environmental Health*, 221:502–509, 2018.

[25] Xinpeng Shen et al. Challenges and opportunities with causal discovery algorithms: Application to alzheimer's pathophysiology. *Scientific Reports*, 10, 2020.

[26] Yan Wang et al. Estimating causal effects of long-term pm 2.5 exposure on mortality in new jersey. *Environmental Health Perspectives*, 124, 2016.

[27] Z Zhang et al. Comparison of micro- and nano-size particle depositions in a human upper airway model. *Journal of Aerosol Science*, 36:211–233, 2005.

[28] JF Gamble. Pm2.5 and mortality in long-term prospective cohort studies: cause-effect or statistical associations? *Environmental Health Perspectives*, 106:535–49, 1998.

[29] C.W.J Granger. Investigating causal relations by econometric models and cross-spectral methods. *Econometrica*, 37:424–438, 1969.

[30] C Hiemstra and J.D Jones. Testing for linear and nonlinear granger causality in the stock price-volume relation. *Journal of Finance*, 49:1639–1664, 1994.

[31] Lei Jiang and Ling Bai. Spatio-temporal characteristics of urban air pollution and their causal relationships: Evidence from beijing and its neighboring cities. *Scientific Reports*, 8, 2018.

[32] S Kanaya. A nonparametric test for stationarity in continuous & time markov processes. *Job Market Paper, University of Oxford*, 2011.

[33] Helmut Lütkepohl. *New Introduction to Multiple Time Series Analysis*. Springer, 2005.

[34] R. G. Norman, M. M. Ahmed, J. A. Walsleben, and D. M. Rapoport. Detection of respiratory events during npsg: nasal cannula/pressure sensor versus thermistor. *Sleep*, 20:1175–1184, 1997.

[35] World Health Organisation. Tackling the global clean air challenge. 2011.

[36] J. Runge. Quantifying information transfer and mediation along causal pathways in complex systems. *Phys. Rev. E*, 92, 2015.

[37] J Runge. Causal network reconstruction from time series: From theoretical assumptions to practical estimation. *Chaos: An Interdisciplinary Journal of Nonlinear Science*, 28, 2018.

[38] J Runge. Conditional independence testing based on a nearest-neighbor estimator of conditional mutual information. *PMLR*, 84:938–947, 2018.

[39] J. Runge, J. Heitzig, N. Marwan, and J.Kurths. Quantifying causal coupling strength: A lag-specific measure for multivariate time series related to transfer entropy. *Phys. Rev. E*, 86, 2012.

[40] A Savitzky and M.J.E Golay. Smoothing and differentiation of data by simplified least squares procedures. *Analytical Chemistry*, 36:1627–39, 1964.

[41] J Schwartz. Air pollution and hospital admissions for the elderly in detroit, michigan. *American Journal of Respiratory and Critical Care Medicine*, 150, 1994.

[42] J Schwartz. What are people dying of on high air pollution days? *Environmental Research*, 64:26–35, 1994.

[43] J Schwartz and MA Bind. Estimating causal effects of local air pollution on daily deaths: effect of low levels. *Environmental Health Perspectives*, 125:23–29, 2017.

[44] Peter Spirtes and Clark Glymour. An algorithm for fast recovery of sparse causal graphs. *Social Science Computer Review*, 9:62–72, 1991.

[45] Peter Spirtes and Kun Zhang. Causal discovery and inference: concepts and recent methodological advances. *Applied Informatics*, 3:1–28, 2016.

[46] Z. Zhang and C. Kleinstreuer. Species heat and mass transfer in a human upper airway model. *International Journal of Heat and Mass Transfer*, 46:4755–4768, 2003.

[47] CM Zigler and F Dominici. Point: Clarifying policy evidence with potential-outcomes thinking - beyond exposure-response estimation in air pollution epidemiology. *American Journal of Epidemiology*, 180, 2014.

[48] D.K. Arvind; Delhi Air Pollution: Health aNd Effects. `http://www.specknet.uk/projects`.

[49] Editorial team; For Those With Asthma, What is Normal Breathing? `https://asthma.net/living/normal-breathing/`.

[50] European Space Agency; Coronavirus: nitrogen dioxide emissions drop over Italy. `http://www.esa.int/ESA_Multimedia/Videos/2020/03/Coronavirus_nitrogen_dioxide_emissions_drop_over_Italy`.

[51] European Space Agency: COVID-19: nitrogen dioxide over China. `https://www.esa.int/Applications/Observing_the_Earth/Copernicus/Sentinel-5P/COVID-19_nitrogen_dioxide_over_China`.

[52] J. Runge; Tigramite. `https://jakobrunge.github.io/tigramite/`.

[53] Python; scikit-learn package. `https://scikit-learn.org/stable/`.

[54] Python; statsmodels package. `https://www.statsmodels.org/stable/index.html`.

[55] S. Palachy; Stationarity in time series. `https://towardsdatascience.com/stationarity-in-time-series-analysis-90c94f27322`.

[56] Wikipedia; Augmented Dickey–Fuller test. `https://en.wikipedia.org/wiki/Augmented_Dickey-Fuller_test`.