Modelling the association between weather and short-term demand for children's intensive care transport services during winter in the South East of England

Samuel Livingstone[a,1,*], Christina Pagel[b], Zejing Shao[a], Elise Randle[c], Padmanabhan Ramnarayan[c]

[a] *Department of Statistical Science, University College London, UK*
[b] *Clinical Operational Research Unit, University College London, UK*
[c] *Children's Acute Transport Service, Great Ormond Street Hospital, London, UK*

Abstract

Data from a paediatric intensive care transport service based in the South East of England between 2006 and 2018 are studied using generalized additive models to investigate the effects of extreme weather on demand in winter. Noticeable increases in daily demand for the service are uncovered after periods of extreme weather, and can be partitioned into two characteristically different phenomena, most pronounced at 2 days and 7 days after a period of particularly low temperature combined with either high or low humidity. The effect is more visible when virus prevalence is accounted for, showing that demand can increase by as much as 30% 7 days after a period of low temperature and low humidity, and 20% 2 days after a period of low temperature and high humidity.

Keywords: Demand forecasting, Paediatric intensive care transport services, Generalized additive models

1. Introduction

Paediatric intensive care transport services operate throughout the United Kingdom as a means of transferring critically-ill children quickly and safely from a local district hospital to a specialist facility. Provision of such services is important as it allows very sick children to receive intensive care level treatment more quickly. Running such a service is not without cost since the services are staffed by highly specialised intensive care clinical teams, however the use of such specialised teams has been shown to increase the survival of these critically-ill children [1]. Retrieval services are quite small (11 locations covering England

*Corresponding author
Email address: samuel.livingstone@ucl.ac.uk (Samuel Livingstone)

and Wales, each with typically 1-2 teams available to transport children) but the population they serve is high risk. In times of high demand, there may not be a team available to retrieve a critically-ill child which could impact on that child's eventual outcome. Being able to predict increases in demand a few days in advance could allow for temporary increases in resourcing to reduce the chance of a team not being available when needed [2].

In this article we analyse 12 years of data (2006-2018) from the Children's Acute Transport Service (CATS), which serves around 50 district general hospitals in the North Thames, Hertfordshire, Bedfordshire, Essex, Norfolk, Suffolk and Cambridgeshire regions of England. CATS, based at Great Ormond Street Hospital, is the paediatric intensive care transport service that performs the most retrievals each year in the UK (e.g. [3]) and one of the largest specialist paediatric retrieval services in Europe [4]. In 2017/18, for example, the service handled nearly 2,400 calls and mobilized a specialist team for more than 1,200 patient transports, an average of 7 advice calls and 3 patient transports per day [3]. The CATS service has been operational since 2001 and since then demand for services has steadily increased year-on-year.

There is a strong seasonal component to demand for CATS services - the number of referrals almost doubles during winter [5, 6]. In addition to modelling these effects, however, a key focus of the present study is short-term weather-induced fluctuations in demand for the service during winter. Members of the CATS team have observed anecdotally that there is a 'cold snap' effect; meaning a temporary spike in referrals shortly after a period of extremely cold weather. The hypothesis is that a period of very low temperature combined with extreme humidity induces an increase in referrals, particularly those related to respiratory illnesses. The connection between weather and respiratory illnesses is well-studied [7, 8]. It is known that viruses tend to be more prevalent in colder temperatures, with many different explanations offered [9]. The effects of humidity appear to vary depending on the type of illness [7], with evidence that both very high and very low humidities have been associated with increasing susceptibility to different types of respiratory condition. We discuss this in more detail in Section 7.1. The CATS population tend to be young children with small airways and vulnerable immune systems who are more likely to develop respiratory distress, so even a small effect could impact demand for services. Therefore, following an incubation period, one would expect an increase in referrals for CATS services. Understanding the size and duration of such an effect to assess the practicality of predicting and flexing the staffing of the service to meet the needs of the children is the main goal of the present work. In Section 7.2 we discuss implications of the study findings for resource planning of the service. We analyse referrals to the service from the period April 2006 to May 2018, which includes both advice calls and demand for transport.

Many previous studies model the association between weather and ambulance demand, but very few with the purpose of analysing short-term changes as a result of periods of extreme weather. Wong & Lai [10] note that including weather variables such as average temperature and relative humidity as covariates in auto-regressive integrated moving-average (ARIMA) models to forecast ambulance demand in Hong Kong decreased the root mean squared error in predictions by 10% for seven-day forecasts and 8.8% for one-day forecasts for the period May 2006 - April 2009. Thornes et al. studied ambulance call-outs in Birmingham [11] from 2007-2011, and note that ambulance call-outs for patients with breathing problems increased significantly in December 2010, the coldest month of the study period.

The link between weather and general hospital admissions has also been studied previously. Lee et al. [12], for example, investigated the association between meteorological factors and visits to a paediatric emergency department in Changwon, Korea from 2005-2014. Using a quasi-Poisson generalized linear model with some non-parametric terms included the authors found that visits increased two days after a rainy or snowy day. The focus of the study, however, was mainly on modelling the *reduction* of visits to hospitals during bad weather, so this effect may well be an artefact of visits being delayed, something that would not be expected of emergency referrals to paediatric intensive care (the focus of the present study). Loh et al. [13] modelled weekly incidence of respiratory infections in a Singapore hospital and associations with both climate and clinical factors. A time series modelling approach (ARIMA) was taken, with linear covariates included for climate variables. It was found that both temperature and relative humidity were negatively correlated with virus incidence. The data, being weekly, was not as granular as that of the present work, which may explain why no lags were needed for the weather variables that were used as model covariates.

2. Materials

The data consists of 12 years of anonymised referrals to the CATS service, from 1 April 2006 until 2 May 2018. There were 26,753 referrals during this period. Each referral record includes a time stamp, the age of the patient, the referring hospital, one of 20 diagnostic categories, and the type of CATS service requested, either an advice call or demand for transport or out of scope. The majority of referrals occur between 09:00-23:00, and the most common diagnostic category is respiratory disease, accounting for 34% of the total. We grouped referrals into into daily (midnight to midnight) counts for modelling purposes. All the data are used in the modelling process, not only the referrals that occurred during the Winter period.

We used two sources of meteorological data that covered the study period: the National Oceanic and Atmospheric Administration (NOAA) Global Daily Historical Climatology Network [14], and weather station data provided by the University of Cambridge Digital Technology Group (DTG) [15]. The NOAA data contains daily temperature readings from multiple weather stations around the world. There are 114 NOAA weather stations in the UK, and 5 of these are within the CATS operating region. Comparisons of data at the 5 stations indicated that temperature did not vary a great deal over the study region, and so the most abundant source of daily data (Heathrow) was used for analysis. The DTG data contained various meteorological recordings, including humidity, made at a weather station in Cambridge, which is geographically at the centre of the CATS region, each at 30 minute intervals throughout the study period. Daily average humidity readings from this source were taken as a proxy for humidity across the CATS operating region.

A summary of the finalised data set used in the analysis is given in Table 1. This also includes virus data, discussed in Section 6.3.

Table 1: The final dataset used for analysis.

Variable	**Source**	**Comment**
Date	CATS	Study day, 1-4415
Day	CATS	Day of the year, 1-365
Calls	CATS	Daily calls made to the CATS operating team
Demand	CATS	Daily requests for transport made to the CATS team
Respiratory calls	CATS	Daily calls related to respiratory illnesses
Respiratory demand	CATS	Daily requests for transport related to respiratory illnesses
Temperature	NOAA	Average temperature (degrees Celsius) on study date, recorded at Heathrow
Humidity	DTG	Average relative humidity (%) on study date, recorded in Cambridge
Virus	PHE	Indicator variable to reflect prevalence of respiratory viruses in England and Wales on study date (see Section 6.3)

CATS: Children's Acute Transport Service, NOAA: National Oceanic and Atmospheric Association, DTG: Digital Technology Group (University of Cambridge), PHE: Public Health England.

3. Exploratory data analysis

To develop a basic understanding of the referrals data we performed a seasonal trend decomposition [16] (Figure 1) using the `R` statistical software [17].

The overall trend in increasing demand for the service over the study period can be clearly seen (using only the respiratory diagnostic category produces a similar result). The large seasonal variation is also striking. To further explore the seasonality a 30-day moving average plot of each individual study year is overlaid on a single graph in Figure 2.

While there is some variation between different years, there is a consistent increase in referrals from November to February, a well known feature of demand for emergency services in health care [2]. Clearly such a seasonal pattern is strongly influenced by meteorological factors, however the focus of the present study is not these seasonal changes, but rather short-term fluctuations

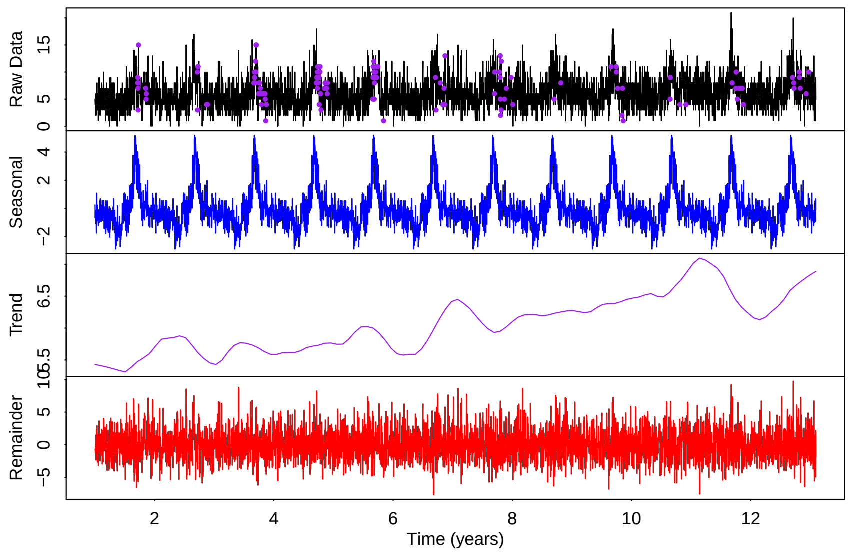

Figure 1: Seasonal trend (STL) decomposition of all referrals across study period. Purple dots in the first plot indicate days in which temperature is in the lowest 10th percentile and humidity is in either the lowest or highest 10th percentile.

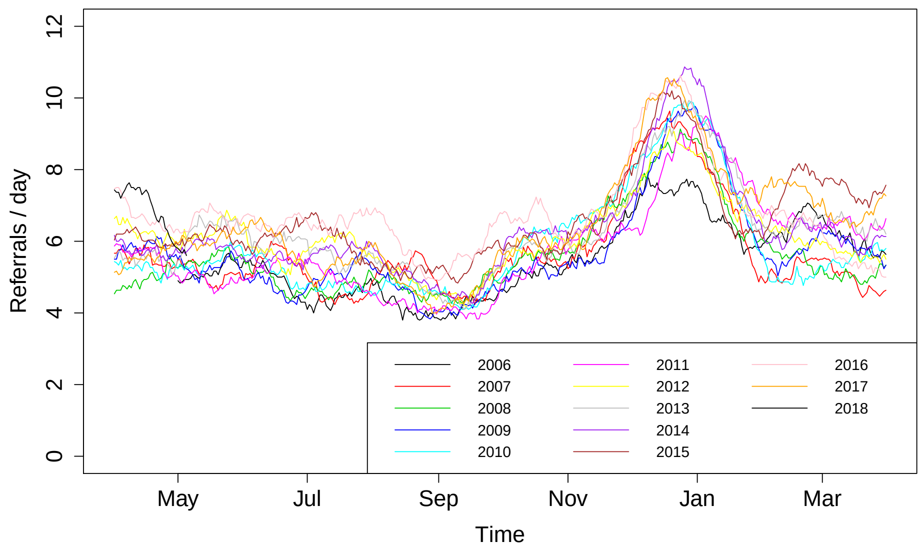

Figure 2: Smoothed referrals/day during each year of the study period, 30-day moving average plot.

in weather patterns and the effects these have on referrals over and above the annual "winter surge" in demand. Controlling for the long-term and seasonal effects is therefore an important modelling component. We note that while there may be short-term fluctuations in demand for other times of year, these are of less importance to the service since there is plenty of capacity to meet extra demand outside of winter (because baseline demand is lower).

Figure 3 shows the average monthly temperature and relative humidity across the study period. This information can be used to assess the degree to which seasonal and short-term weather effects can be separated. The figure shows that average temperatures in the coldest months are 5.6 degrees Celsius, and in the warmest month 18.7 degrees. Similarly, average relative humidity ranges from 70-90%. The two covariates are clearly negatively associated. The 10th and 20th percentiles for temperature are 2.44 and 5.13 degrees Celsius, and for relative humidity the 10th and 20th percentiles are 65.5% and 70.0% and the 80th and 90th are 89.7% and 93.6%.

4. Modelling approach

Aggregating the data into daily counts provides a straightforward framework for probabilistic modelling. The natural starting point is to assume that the number of referrals/advice calls/requests for transport per day will approximately follow a Poisson distribution. If Y_t represents some appropriate response

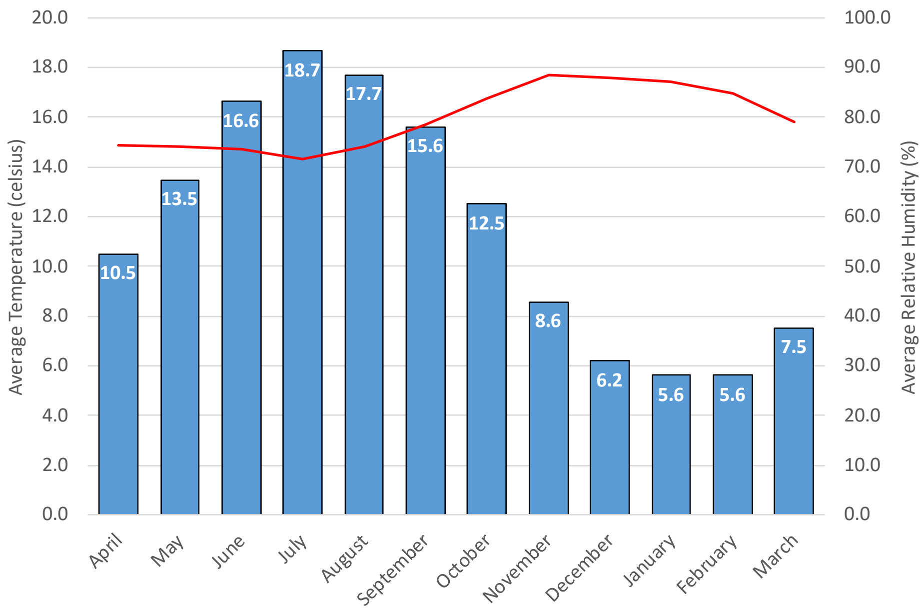

Figure 3: Average temperature (blue bars) and relative humidity (red line) values in each calendar month, taken over the study period April 2006 - April 2018, using the NOAA and DTG datasets.

(e.g. total number of referrals) on day t, then we assume $Y_t \sim \text{Poisson}(\mu_t)$. The mean μ_t can be chosen to depend on a number of covariates, and ideally this is done in such a way that useful inferences can be made about these. We chose the generalized additive modelling framework to construct an appropriate model for μ_t, details of which are given in the next subsection.

The daily counts for the CATS service are small and building a model for Poisson data with relatively low daily counts is not straightforward. The coefficient of variation for Y_t is $1/\sqrt{\mu_t}$, meaning that when counts are small, as is the case here, the level of relative variation in the data is higher. As a result, building a good predictive model is a challenge. Aggregating the data further into weekly counts would alleviate the issue, but also make it impossible to capture the short-term effects of interest in this study. Nonetheless, for inference purposes an informative model at the daily level can still be built. In Section 6 we explore various extensions to the Poisson approach, such as models with auto-correlated residuals, controlling for over- or under-dispersion, and assuming different distributions for the response.

4.1. Generalized additive models

Generalized additive models (GAMs) [18] are a flexible approach to modelling data in which different variables may be related through some unknown nonlinear mechanism. Given data $\{(x_{1t}, ..., x_{pt}, y_t)\}_{t=1}^n$, we suppose that Y_t follows some known exponential family distribution (e.g. Normal, Poisson, Binomial) with $E[Y_t] = \mu_t$, and seek to fit a model of the form

$$g(\mu_t) = \alpha + \sum_{j=1}^{p} f_j(x_{jt}), \tag{1}$$

where $\alpha \in \mathbb{R}$, each f_j is a smooth real-valued function, and g is called the link function. One can also include terms of the form $f_{jk}(x_{jt}, x_{kt})$ in the above expression, to allow for interactions between different covariates. GAMs offer an attractive compromise between black-box predictive modelling approaches that can be difficult to interpret and more restrictive linear modelling frameworks that can often impose unrealistic assumptions.

The functions f_j and the intercept term are inferred from the data in a GAM: the user need only specify the distribution for the response Y_t and the form of the link function g. In the case of $Y_t \sim \text{Poisson}(\mu_t)$, then the canonical choice of link function is $g(\mu_t) = \log(\mu_t)$, which ensures that μ_t is positive. To avoid issues of over-fitting, the model is typically trained using penalized maximum likelihood methods, with the level of penalisation determined either by generalized cross-validation or empirical Bayes techniques [18, Ch. 4]. The individual f_j's are learned by first specifying a set of basis functions b_{jk} for $1 \le k \le k_j$, and then using penalised least squares with a degree of smoothing/regularisation controlled by some parameter $\lambda_j \in \mathbb{R}$ to estimate the parameters γ_{jk} in the specification $f_j(x) = \sum_k \gamma_{jk} b_{jk}(x)$. A spline basis using k_j knots within the range of the data is typically used [18, Ch. 4]. There are therefore two mechanisms by which the degree of smoothness of each f_j can be controlled to prevent overfitting: choosing a smaller number of knots (controlled by the user), or allowing a larger value for λ_j (which is typically tuned to minimise out-of-sample error during the inference procedure).

We used the `mgcv` package [19] within the `R` statistical software [17] to fit models to the data. We chose the restricted maximum likelihood approach to estimating the degree of smoothing parameter λ_j of each smooth term f_j, which generally gives more stable results in the presence of model miss-specification where generalized cross-validation can lead to under-smoothing (see [20, 21] for more on this).

4.2. Model building

We fitted several models to the referrals and meteorological data, using different responses and different additive structures for $\log(\mu_t)$. Each formulation included features designed to capture both the long-term and seasonal behaviour of the response over the study period, as well as the short-term weather effects of our specific interest.

A simple starting point is the collection of lag$-k$ models of the form

$$\begin{aligned} \log(\mu_t) =& \alpha + f_1(t) + f_2(\text{day}_t) + f_3(\text{temp}_{t-k}, \text{humid}_{t-k}) \\ &+ \sum_{j=1}^{6} \beta_j \mathbb{I}(\text{dow}_t = j) + \gamma h_t, \end{aligned} \tag{2}$$

where t denotes the day of the study (from $1-4415$), day_t denotes the day of the year on study day t (from $1-365$, where 1 represents 1st January), and the variables temp_i and humid_i represent the temperature and humidity on study day i. The integer k indexes a particular model, and determines the lag, or in real terms the assumed period of time required for any medical conditions resulting from adverse weather to manifest. The variable $h_t = 1$ if study day t was a public holiday in England, and 0 otherwise, and $\text{dow}_t \in \{1, ..., 6\}$ is a factor to control for the day of the week. The indicator function $\mathbb{I}(x) := 1$ if the Boolean variable x is true and 0 otherwise.

Some variations of the above formulation were also tested, such as removing f_2 and allowing f_1 to capture both long-term and seasonal effects. This approach has the advantage of being able to capture different seasonal effects in different years, at the expense of requiring many more basis functions to estimate the function f_1 appropriately, which significantly increases the effective number of parameters to be inferred. A risk in this formulation is identifiability issues between temporal and weather effects, since the association between the time of year and temperature is high. The use of the f_1, f_2 formulation shown in (2) addresses these issues to some degree, as a high amount of smoothing can be imposed on f_1 to ensure that only an overall trend is captured, and f_2 is restricted to fit an average seasonal affect across multiple years, meaning short-term fluctuations resulting from colder than usual temperatures are smoothed out and would contribute only to f_3. Cyclic cubic splines were used in the estimation of f_2 to ensure that there were no discontinuities between the start and end of each year (see Figure 4 for an example).

It should be noted that for any particular choice of k (2) is almost certainly miss-specified, as it is very unlikely that a single lag will capture the heterogeneity in lag times experienced by different patients after an extreme weather event. It does, however, serve as a starting point for analysis to give an approximate picture of the impact of any single value of k. To give a more accurate picture of the impact of different lags a *single index* model formulation was then used, (e.g. [18, Ch. 7]), which has the formulation

$$\begin{aligned}\log(\mu_t) =& \alpha + f_1(t) + f_2(\text{day}_t) + f_3\left(\sum_{k=0}^{K} w_k \text{temp}_{t-k}, \sum_{k=0}^{K} w_k \text{humid}_{t-k}\right) \\ &+ \sum_{j=1}^{6} \beta_j \mathbb{I}(\text{dow}_t = j) + \gamma h_t,\end{aligned} \tag{3}$$

with the weights w_k learned from the data, and satisfying $w_k > 0$ and $\sum_{k=1}^{K} w_k = 1$. The single index approach allows weather effects at multiple time lags to be incorporated into the same model, without overburdening the modelling task by attempting to estimate several bi-variate smooth terms simultaneously for different lags.

The single index and fixed lag models can be combined to provide a good understanding of the overall short-term effects of weather on referrals, as the weights in the former reveal the most influential lag terms, and the individual lag formulations provide a more nuanced (albeit approximate) understanding of the effects at each individual lag. We illustrate how this combination leads to a particularly expressive way of capturing the effects of weather on referrals in the next section.

5. Results

All model formulations were tested using both all referrals and demand for transport only as the response, and in addition either across all diagnostic categories or focusing exclusively on respiratory referrals. The results across the different formulations were in fact quite similar, providing some evidence that the identified effects are real. The assumption is that referrals focused on respiratory conditions are the most likely to be affected by weather fluctuations (and respiratory conditions drive the annual winter surge), but equally restricting to a smaller data set in this way increases the coefficient of variation, making it harder to estimate the different effects reliably. The results presented in the below sections are using all referrals as the response, which proved the most robust.

5.1. Fixed lag models

Generally the fixed lag models varied very little across different choices of k in terms of inferring the overall trend $f_1(t)$ and the seasonal pattern $f_2(\text{day}_t)$. An example in the case $k = 4$ is given in Figure 4. The overall increasing trend and increase in demand during the Winter period can be clearly seen in the forms of f_1 and f_2.

Contour plots of the estimated bi-variate smooth function of temperature and humidity (f_3) for fixed lags from $k = 0$ up to $k = 10$ are shown in Figure 5. Initially with no lag there appears to be little association between the weather variables and demand for the service, which makes sense as extreme weather is unlikely to make children critically-ill instantaneously. The effects between $k = 1$ and $k = 5$, most pronounced for $k = 2$, suggest that a combination of low temperature and high humidity a few days previously results, on average, in an increase in demand of up to 20% depending on the severity of the weather. This is consistent with the 'cold snap' hypothesis described in the introduction. When $k \geq 6$, however, a noticeably different pattern can be observed in the plots. Here high humidity does not appear to be important, and if anything lower humidity (combined with low temperature) seems to be an indicator of an increase in referrals. The models $k = 7$ and $k = 10$ in particular score well according to the AIC (Akaike Information Criterion [22]) and REML (restricted maximum likelihood) scores of model fit (REML scores shown in Figure 6).

5.2. Single index model

We have seen above that various single lag models capture different levels and types of association between weather and referrals. The single index formulation allows us to infer which of these lags are the most informative in predicting

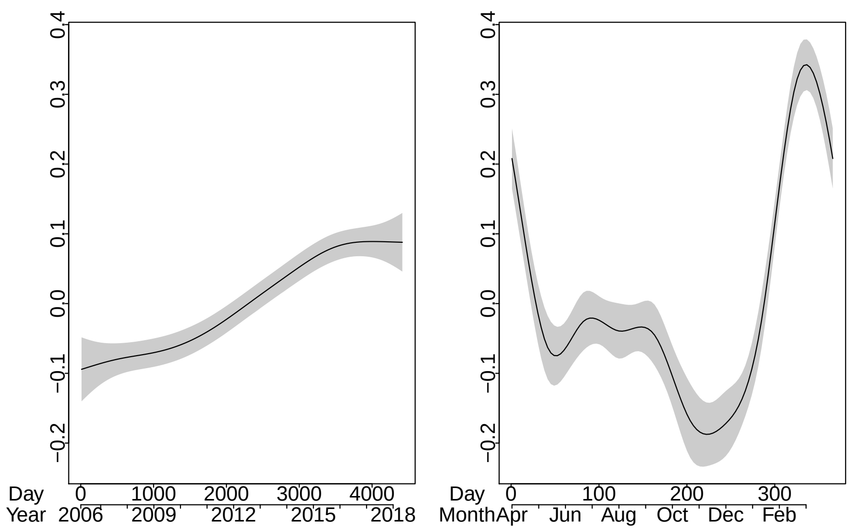

Figure 4: Estimated smooth effects of time (left) and day of the year (right) on the logged-mean total number of calls, taken from lag-4 model.

referrals, and also to some degree control for any interactions between different lags by including them in the same model. Figure 7 shows the estimated values of each weight w_k for $k \in \{0, ..., 10\}$.

The single index model attaches the most prominent weights to the 2-day and 7-day lag formulations. This is intuitive, as the 2-day and 7-day lag in particular seem to represent quite different effects. The 2-day lag model is also the best REML fit among the $k = 0$ to 5 suite of models, all of which appear to estimate characteristically similar weather effects as evidenced by Figure 5. Similarly the $k = 7$ model is among the best REML fits from the models with $6 \leq k \leq 9$ lags. The $k = 10$ model may again represent a slightly different effect.

Full estimates for each of the single index model weights, together with estimates for the parameters $\beta_1, ..., \beta_6$, α and γ are given in Appendix B.

6. Extensions

6.1. Residual auto-correlation

The time series structure of the referrals data suggests that there might be auto-correlation in the model residuals. Some example residual auto-correlation plots are shown in Figure A.11 in Appendix A, showing that this residual auto-correlation did not seem to be a feature of our fitted models. One possible

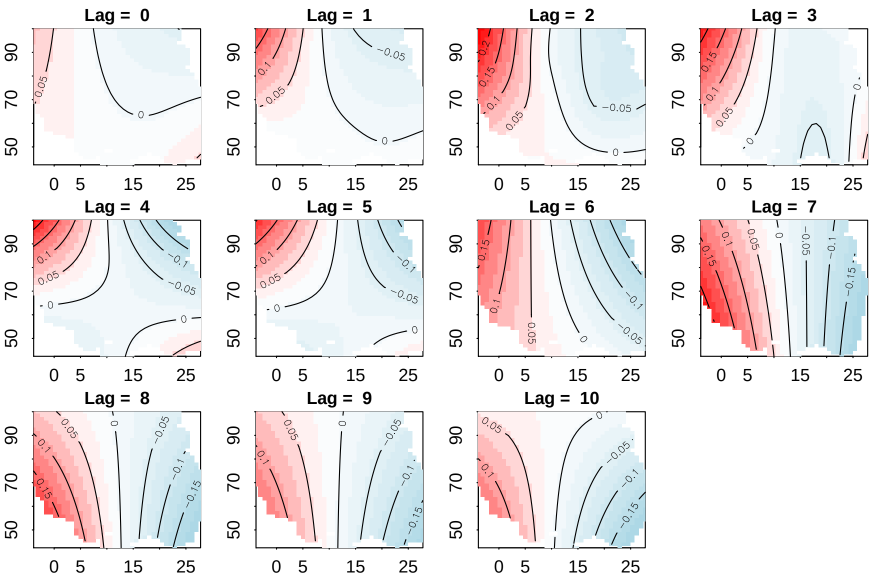

Figure 5: Bi-variate smooth term plots for each of the fixed lag models. All x-axes are temperature, all y-axes are humidity. Red indicates higher demand and blue lower demand.

explanation is that the data are counts with relatively small numbers and hence large relative variation, which may dominate any residual auto-correlation effects that are present. Another is that the temporal structure is fully captured using the functions f_1 and f_2. Nonetheless we did fit a Generalized Additive Mixed Model to the data (GAMM) [18, Ch. 7] with an AR(1) residual structure to test for any relevant patterns. These models are a challenge to fit using the current suite of implementable methods in the `mgcv` package, as they are computationally intensive and not as numerically robust as GAMs, but in the single lag $k = 4$ case the model fitting process converged successfully with an estimated value for the AR(1) correlation parameter $\hat{\rho} = 0.05$, suggesting that very little auto-correlation was present in the residuals. The estimated bi-variate smooth function of temperature and humidity in the AR(1) residual case is shown in Figure A.10 of Appendix A, and is consistent with the uncorrelated case.

6.2. Different response distributions

Over-dispersion is common in regression models for count data in which the response is assumed to be Poisson, as the imposed structure implies that $\text{Var}(Y_t) = \mu_t$. Often some desired covariates cannot be included in the model, which can have the effect of inflating the variance of Y_t. One popular solution is the *quasi-Poisson* approach [23]. Rather than specify an entire distribution for Y_t, in the quasi-Poisson formulation only the mean μ_t and variance are directly modelled, and the variance is assumed to be $\phi\mu_t$, where $\phi > 0$ is called the

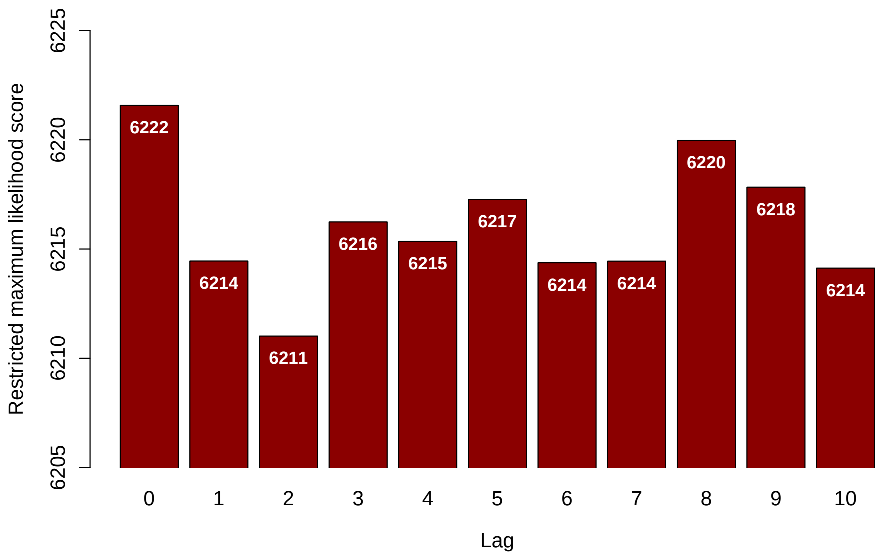

Figure 6: Restricted maximum likelihood values for each of the fixed lag models (a lower score corresponds to a better fitting model).

dispersion parameter. An estimate $\hat{\phi} \gg 1$ indicates significant over-dispersion. To test for over-dispersion quasi-Poisson versions of each of the above fixed lag models were fitted to the referrals data. In every case $\hat{\phi} < 1.04$, suggesting that very little over-dispersion was present, and the estimates for f_1, f_2 and f_3 were essentially unchanged. Using an alternative method for assessing over-dispersion, assuming that Y_t instead follows a negative Binomial distribution, also showed minimal over-dispersion.

There are more days with zero demand in the data than would be expected under the Poisson assumption. A natural way to capture such effects is using a zero-inflated model, in which the response is assumed to be zero with some probability $(1-p)$, and conditional on not being zero it is modelled as a truncated Poisson variable (truncated to be ≥ 1, e.g. [24]) . Zero-inflated models can allow for a much more flexible model structure, but can also sometimes be ineffective, particularly when the covariates do appear to have a relationship with whether or not the response takes a zero value. We fitted a zero-inflated Poisson GAM using the `mgcv` package, which offers functionality to estimate a fixed p from the data, and found $\hat{p} = 0.98$, indicating that the best estimates from the data and proposed modelling structure were that only 2% of the data should be zero. This contrasted significantly with the reality in which 18% of total referrals were zero. The likely explanation is that the covariates are indeed associated with the response being equal to or greater than zero. Unsurprisingly given the

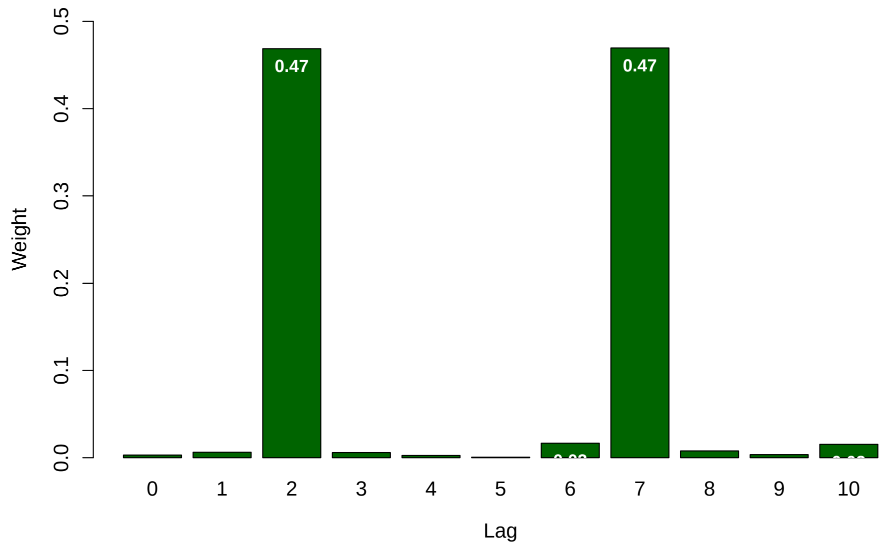

Figure 7: Weights associated with each lag in the fitted single index model.

estimated value of p, the zero-inflated Poisson models gave very similar results to the Poisson models in terms of inferences for all other effects.

6.3. Including virus information

Adverse short-term weather is unlikely to result in high referrals in isolation, particularly for respiratory referrals. The large majority of sick children will have been exposed to some pathogen such as influenza or the common cold virus. To capture this 'probability of exposure' effect some indication of virus prevalence was incorporated into the model. The data we used for this purpose are the routine respiratory infection data weekly reports provided by Public Health England (e.g. [25]). The reports provide information about the prominence of various respiratory infections in England and Wales on a weekly basis. Data are available throughout the study period, however only at a weekly level and with 7% of weeks missing. Because of the coarse granularity and missingness we incorporated this information into the model using a simple indicator variable

$$\text{virus}_t := \mathbb{I}(A_t > q_A \text{ or } B_t > q_B), \tag{4}$$

where A_t and B_t denote the number of recorded reports of influenza A and B infection in England and Wales during the week of study day t, q_A and q_B denote the 90% quantiles for weekly reports of Influenza A and B respectively across the study period, and $\mathbb{I}(x) = 1$ if x is true and 0 otherwise. In words, the virus variable is a factor with two categories, category zero to indicate lack of outbreak of influenza, and category one to indicate that the number of reports of some form of influenza virus are unusually high.

The new extended lag-k model formulation is identical to (2), except that the bi-variate smooth term is replaced with the term

$$\sum_{i=1}^{2} f_{2+i}(\text{temp}_{t-k}, \text{humid}_{t-k})\mathbb{I}(\text{virus}_{t-k} = i - 1). \tag{5}$$

In effect, the model now fits two bi-variate smooth functions at each lag, one to assess the effects of temperature and humidity on days when the virus indicator is zero, and another when it is one (implying that the viral load is high).

The addition of the virus covariate had a small but noticeable effect on the model fitting. The functions f_1 and f_2 were essentially unchanged, but the same is not true of the bi-variate smooth terms, shown in Figure 8 for the 2 day, 4 day, 7 day and 10 day lag models.

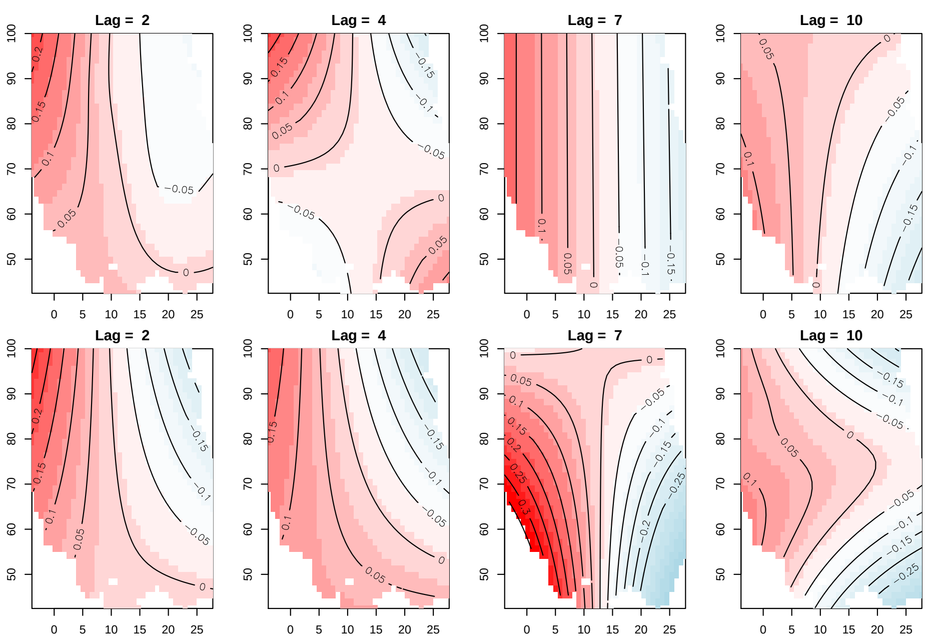

Figure 8: Bi-variate smooth term plots for the fixed lag 2, 4, 7 and 10 models, with virus indicator low (first row) and high (second row). Red indicates high demand and blue lower demand.

In each case there is some refinement to the inference. In particular in the lag−2 and lag−7 cases the effect size is increased in the case of high viral load, which can now be as high as 30% for the 7 day lag model. Wald-type tests to assess the significance of smooth terms in generalized additive models developed in [26] based on the approach of [27] indicated that both smooth terms included in the virus models were significant ($p < 0.01$), but an analysis of deviance comparing the lag-2, 4 and 7 models with and without the virus terms was less conclusive ($p = 0.99$ for lag-2, $p = 0.34$ for lag-4, $p = 0.11$ for lag-7, $p = 0.42$ for lag-10). The apparent disagreement between these conclusions can be explained by the fact the Wald tests are assessing whether or not it is plausible that the true smooth term could be the zero function. The low p-value here indicates that this is not very plausible. The deviance measure, however, assesses the predictive ability of the model, and whether it is worth including this extra term in terms of predictive ability as measured by the deviance, which is a sum across all observations in the dataset. This score is only slightly affected because there are relatively few days in the dataset in which temperature is low, humidity is high/low and in addition the viral load is high. As such, this measure of predictive ability is not influenced much by the predictions for these days being slightly better, as they are a small fraction of the total number of study days from which the deviance is calculated.

7. Discussion

In this paper we have established for the first time a similar association between periods of low temperature and high humidity and increased demand for specialised paediatric ambulance services. From both fixed lag models and the single index model we see an increase in demand for transport to intensive care at 2 days (low temperature, high humidity) and 7 days (low temperature, low humidity) following extreme weather.

7.1. Clinical plausibility

The findings from this analysis are clinically plausible, with the two different effects explainable by different pathologies. Anecdotally the CATS team have recognised that after a period of cold weather there is an increase in children referred who have an underlying reactive airway disease such as asthma. Many of these children report changes in weather as triggers for their reactive airway disease. This would be consistent with our findings here that a drop in temperature and associated high humidity may increase the number of children being referred for transport to a paediatric intensive care unit (PICU) 2 days later. Several studies have noted an association between high humidity and onset of respiratory conditions such as asthma [28, 29], with one possible causal mechanism an increased prevalence of allergenic mould and dust mites in households [30, 31, 32].

The children presenting at day 7 following a drop in temperature but this time a low humidity would be consistent with infective respiratory presentations. It is known that in low temperature and humidity conditions many viruses can survive for longer on surfaces [33, 34, 35]. Other causal mechanisms for virus spread in low temperatures have also been suggested, such as crowding in indoor spaces [9] and a cooling of the nasal cavity rendering the immune response less effective [36]. In this modelling, the additional effect of high viral load in the community increased demand for transport by approximately 30% at 7 days after the weather change.

7.2. Resource planning

Studies exploring the association of ambulance call outs [11, 37, 38] and emergency departments [39] with extreme weather events have found significant increases in demand during extreme cold or extreme hot periods for adults. They suggest that such associations could be used to flexibly increase resources in advance of extreme weather warnings to improve access to critical services. Papadakis et al. [37] for instance provided illustrative tables of expected potential increased demand during a 7-day cold snap that could be used for service planning.

Services for critical care transport from local hospitals to paediatric intensive care units are highly specialised. While overall demand is low compared with adult services or non-specialised ambulance services, they are resource intensive and care for the sickest children. With only 11 locations and a few teams within each location serving England and Wales [40], there is potential for even small increases in demand to stretch capacity – particular in winter when services are already stretched [5]. Previous work has shown how the provision of an extra team during day times only at the largest retrieval service during winter was associated with fewer refusals due to no available team [2]. The association found in this work has the potential for improving allocation of additional resource across the country. As part of the DEPICT study on paediatric critical care retrieval [41], an optimisation framework has been developed that can be used to allocate different numbers of teams to different retrieval locations depending on season and time of day to maximise the availability of teams and minimise time from child referral to the team arriving at the bedside [42]. In that work it was shown that the addition of even one or two teams in key locations could significantly reduce times to bedside. The study considered adding teams for the whole of a given season over a specific time, however. In future work we aim to use the optimisation framework to dynamically allocate extra teams to locations depending on local weather forecasts and knowledge of existing demand (for instance, if current demand is low an extra team might not be needed even in a period of extreme weather).

The next steps for this work are to first improve and validate the models using national data and then to integrate the updated models into a software tool that could be used by local services that combines the models, the developed optimisation framework, local weather forecasts and recent demand to suggest periods during which additional resource would be beneficial.

7.3. Limitations

A limitation of the current study is that the viral data used was for influenza A and B. These are viral infections that cause many adults and some children to be admitted to intensive care but in children they are not the most common viral pathogens causing PICU admission each winter, the most common being respiratory syncytial virus (RSV). It may be that conditions allowing the rise of influenza A and B are also optimal for RSV propagation but ideally data for RSV could be used.

As highlighted by a reviewer of the manuscript, it is also possible that other diseases may increase in incidence in response to weather changes. Parsons et al. [43] study associations between weather and traumatic injury. Although hot weather appears to be a stronger predictor than cold, a 3.2% increase in occurrence for each 5 degree (Celsius) drop in minimum temperature is observed among adults, and a 7.9% increase is observed in the presence of snow. The same pattern is not, however, present among children. There is significantly more literature connecting heart disease and cold weather. Donaldson and Keatinge [44] found that extended periods of cold temperature below 15 degrees Celsius are often followed rapidly by an increase in heart disease mortalities. Anderson and Riche [45] argue that much of this increased incidence is, however, a consequence of an increase in incidence of respiratory conditions. Pell and Cobbe [46] note that the seasonal variation in heart disease incidence varies across age groups, so further work is needed to understand specific effects among children. There is also evidence that cerebrovascular disease incidence and mortality increases in Winter as a result of colder weather [47, 48]. Again the precise effects among children are less well understood. There is also some evidence to suggest that epilepsy [49] and cluster headaches [50] can be influenced by cold weather. Understanding in deeper granularity the nature of demand increases is certainly an interesting area of future work.

Acknowledgment

The authors thank Ziyi Wang for preliminary investigations as part of her MSc project. SL thanks Matteo Fasiolo and Giampiero Marra for useful discussions.

References

[1] P. Ramnarayan, K. Thiru, R. C. Parslow, D. A. Harrison, E. S. Draper, K. M. Rowan, Effect of specialist retrieval teams on outcomes in children admitted to paediatric intensive care units in england and wales: a retrospective cohort study, The Lancet 376 (9742) (2010) 698–704.

[2] C. Pagel, D. Lutman, E. Polke, S. Ray, P. Ramnarayan, Managing the winter surge in demand for resources, British Journal of Healthcare Management 22 (7) (2016) 370–379.

[3] P. I. C. Society, Paediatric Intensive Care Audit Network Annual Report 2018: Tables and Figures, Tech. rep., University of Leeds and University of Leicester, accessed: 2020-05-19 (2018).
URL `https://www.picanet.org.uk/wp-content/uploads/sites/25/2018/12/PICANet_2018_Annual_Report_Tables_and_Figures_v3.0-compressed.pdf`

[4] Children's Acute Transport Service, http://site.cats.nhs.uk/, accessed: 2020-05-19.

[5] C. Pagel, P. Ramnarayan, S. Ray, M. J. Peters, A novel method to identify the start and end of the winter surge in demand for pediatric intensive care in real time, Pediatric Critical Care Medicine 16 (9) (2015) 821–827.

[6] C. Pagel, P. Ramnarayan, S. Ray, M. J. Peters, Development and implementation of a real time statistical control method to identify the start and end of the winter surge in demand for paediatric intensive care, European Journal of Operational Research 264 (3) (2018) 847–858.

[7] J. W. Tang, The effect of environmental parameters on the survival of airborne infectious agents, Journal of the Royal Society Interface 6 (suppl_6) (2009) S737–S746.

[8] J. D. Tamerius, J. Shaman, W. J. Alonso, K. Bloom-Feshbach, C. K. Uejio, A. Comrie, C. Viboud, Environmental predictors of seasonal influenza epidemics across temperate and tropical climates, PLoS pathogens 9 (3) (2013) e1003194.

[9] E. Lofgren, N. H. Fefferman, Y. N. Naumov, J. Gorski, E. N. Naumova, Influenza seasonality: underlying causes and modeling theories, Journal of virology 81 (11) (2007) 5429–5436.

[10] H.-T. Wong, P.-C. Lai, Weather factors in the short-term forecasting of daily ambulance calls, International journal of biometeorology 58 (5) (2014) 669–678.

[11] J. E. Thornes, P. A. Fisher, T. Rayment-Bishop, C. Smith, Ambulance call-outs and response times in birmingham and the impact of extreme weather and climate change, Emergency Medicine Journal 31 (3) (2014) 220–228.

[12] H. J. Lee, M. H. Jin, J. H. Lee, The association of weather on pediatric emergency department visits in changwon, korea (2005–2014), Science of The Total Environment 551 (2016) 699–705.

[13] T. Loh, F. Lai, E. Tan, K. Thoon, N. Tee, J. Cutter, J. Tang, Correlations between clinical illness, respiratory virus infections and climate factors in a tropical paediatric population, Epidemiology & Infection 139 (12) (2011) 1884–1894.

[14] Global Historical Climatology Network, https://www.ncdc.noaa.gov/data-access/land-based-station-data/land-based-datasets/global-historical-climatology-network-ghcn, accessed: 2020-05-19.

[15] University of Cambridge Digital Technology Group Weather Data, https://www.cl.cam.ac.uk/research/dtg/weather/, accessed: 2020-05-19.

[16] R. B. Cleveland, W. S. Cleveland, J. E. McRae, I. Terpenning, Stl: A seasonal-trend decomposition, Journal of official statistics 6 (1) (1990) 3–73.

[17] R Core Team, R: A Language and Environment for Statistical Computing, R Foundation for Statistical Computing, Vienna, Austria (2020). URL https://www.R-project.org/

[18] S. N. Wood, Generalized additive models: an introduction with R, CRC press, 2017.

[19] S. Wood, M. S. Wood, Package 'mgcv', R package version 1 (2015) 29.

[20] S. N. Wood, Fast stable restricted maximum likelihood and marginal likelihood estimation of semiparametric generalized linear models, Journal of the Royal Statistical Society: Series B (Statistical Methodology) 73 (1) (2011) 3–36.

[21] P. T. Reiss, R. Todd Ogden, Smoothing parameter selection for a class of semiparametric linear models, Journal of the Royal Statistical Society: Series B (Statistical Methodology) 71 (2) (2009) 505–523.

[22] H. Akaike, Information theory and an extension of the maximum likelihood principle, in: Selected papers of hirotugu akaike, Springer, 1998, pp. 199–213.

[23] R. W. Wedderburn, Quasi-likelihood functions, generalized linear models, and the gauss—newton method, Biometrika 61 (3) (1974) 439–447.

[24] A. F. Zuur, E. N. Ieno, N. J. Walker, A. A. Saveliev, G. M. Smith, Zero-truncated and zero-inflated models for count data, in: Mixed effects models and extensions in ecology with R, Springer, 2009, pp. 261–293.

[25] Respiratory infections: laboratory reports 2015, https://www.gov.uk/government/publications/respiratory-infections-laboratory-reports-2015, accessed: 2020-05-19.

[26] S. N. Wood, On p-values for smooth components of an extended generalized additive model, Biometrika 100 (1) (2013) 221–228.

[27] G. Marra, S. N. Wood, Coverage properties of confidence intervals for generalized additive model components, Scandinavian Journal of Statistics 39 (1) (2012) 53–74.

[28] A. V. Arundel, E. M. Sterling, J. H. Biggin, T. D. Sterling, Indirect health effects of relative humidity in indoor environments., Environmental health perspectives 65 (1986) 351–361.

[29] S. Weiland, A. Hüsing, D. Strachan, P. Rzehak, N. Pearce, Climate and the prevalence of symptoms of asthma, allergic rhinitis, and atopic eczema in children, Occupational and environmental medicine 61 (7) (2004) 609–615.

[30] D. Strachan, C. Sanders, Damp housing and childhood asthma; respiratory effects of indoor air temperature and relative humidity., Journal of Epidemiology & Community Health 43 (1) (1989) 7–14.

[31] R. E. Dales, H. Zwanenburg, R. Burnett, C. A. Franklin, Respiratory health effects of home dampness and molds among canadian children, American journal of epidemiology 134 (2) (1991) 196–203.

[32] J. K. Peat, J. Dickerson, J. Li, Effects of damp and mould in the home on respiratory health: a review of the literature, Allergy 53 (2) (1998) 120–128.

[33] L. M. Casanova, S. Jeon, W. A. Rutala, D. J. Weber, M. D. Sobsey, Effects of air temperature and relative humidity on coronavirus survival on surfaces, Applied and environmental microbiology 76 (9) (2010) 2712–2717.

[34] A. C. Lowen, S. Mubareka, J. Steel, P. Palese, Influenza virus transmission is dependent on relative humidity and temperature, PLoS pathogens 3 (10) (2007) e151.

[35] A. C. Lowen, J. Steel, Roles of humidity and temperature in shaping influenza seasonality, Journal of virology 88 (14) (2014) 7692–7695.

[36] R. Eccles, An explanation for the seasonality of acute upper respiratory tract viral infections, Acta oto-laryngologica 122 (2) (2002) 183–191.

[37] G. Papadakis, Z. Chalabi, J. E. Thornes, Ambulance service resource planning for extreme temperatures: Analysis of ambulance 999 calls during episodes of extreme temperature in london, uk, Atmosphere 9 (5) (2018) 182.

[38] E. Andrew, Z. Nehme, S. Bernard, M. J. Abramson, E. Newbigin, B. Piper, J. Dunlop, P. Holman, K. Smith, Stormy weather: a retrospective analysis of demand for emergency medical services during epidemic thunderstorm asthma, bmj 359.

[39] H. Hughes, R. Morbey, T. Hughes, T. Locker, T. Shannon, C. Carmichael, V. Murray, S. Ibbotson, M. Catchpole, B. McCloskey, G. Smith, A. Elliot, Using an emergency department syndromic surveillance system to investigate the impact of extreme cold weather events, Public Health 128 (7) (2014) 628–635.

[40] M. King, P. Ramnarayan, S. E. Seaton, C. Pagel, Modelling the allocation of paediatric intensive care retrieval teams in england and wales, Archives of disease in childhood 104 (10) (2019) 962–966.

[41] P. Ramnarayan, R. Evans, E. S. Draper, S. E. Seaton, J. Wray, S. Morris, C. Pagel, Differences in access to emergency paediatric intensive care and care during transport (depict): study protocol for a mixed methods study, BMJ open 9 (7) (2019) e028000.

[42] E. Kung, S. E. Seaton, P. Ramnarayan, C. Pagel, Using a genetic algorithm to solve a non-linear location allocation problem for specialised children's ambulances in england and wales, Health Systems (2021) 1–11.

[43] N. Parsons, M. Odumenya, A. Edwards, F. Lecky, G. Pattison, Modelling the effects of the weather on admissions to UK trauma units: a cross-sectional study, Emergency medicine journal 28 (10) (2011) 851–855.

[44] G. C. Donaldson, W. R. Keatinge, Early increases in ischaemic heart disease mortality dissociated from and later changes associated with respiratory mortality after cold weather in south east England, Journal of Epidemiology & Community Health 51 (6) (1997) 643–648.

[45] T. Anderson, W. Le Riche, Cold weather and myocardial infarction, The Lancet 295 (7641) (1970) 291–296.

[46] J. Pell, S. Cobbe, Seasonal variations in coronary heart disease, Qjm 92 (12) (1999) 689–696.

[47] The Eurowinter Group, Cold exposure and winter mortality from ischaemic heart disease, cerebrovascular disease, respiratory disease, and all causes in warm and cold regions of Europe, The Lancet 349 (9062) (1997) 1341–1346.

[48] Y. Zhang, S. Li, X. Pan, S. Tong, J. J. Jaakkola, A. Gasparrini, Y. Guo, S. Wang, The effects of ambient temperature on cerebrovascular mortality: an epidemiologic study in four climatic zones in china, Environmental Health 13 (1) (2014) 1–12.

[49] K.-L. Chiang, J.-Y. Lee, Y.-M. Chang, F.-C. Kuo, C.-Y. Huang, The effect of weather, air pollution and seasonality on the number of patient visits for epileptic seizures: A population-based time-series study, Epilepsy & Behavior 115 (2021) 107487.

[50] Y.-J. Lee, Y.-T. Chen, S.-M. Ou, S.-Y. Li, A. C. Yang, C.-H. Tang, S.-J. Wang, Temperature variation and the incidence of cluster headache periods: A nationwide population study, Cephalalgia 34 (9) (2014) 656–663.

Appendix A. Additional figures

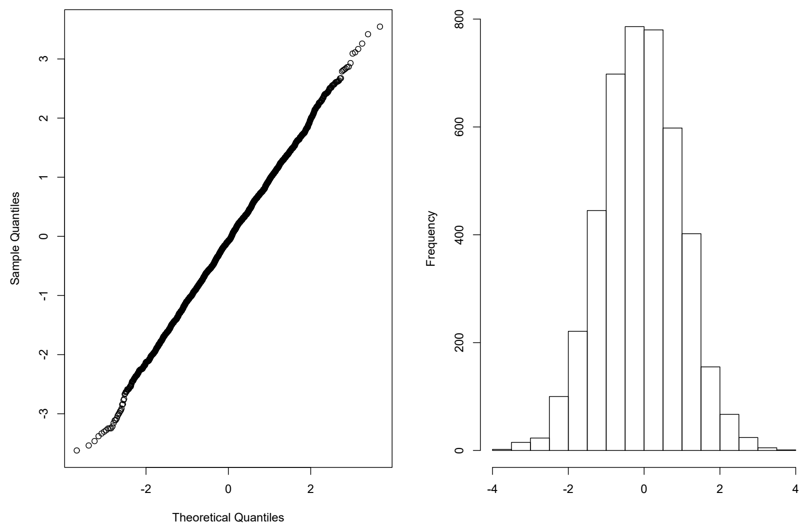

Figure A.9: Example QQ-plot against the standard Normal distribution and histogram of deviance residuals, taken from the lag-4 model.

Appendix B. Parameter estimates

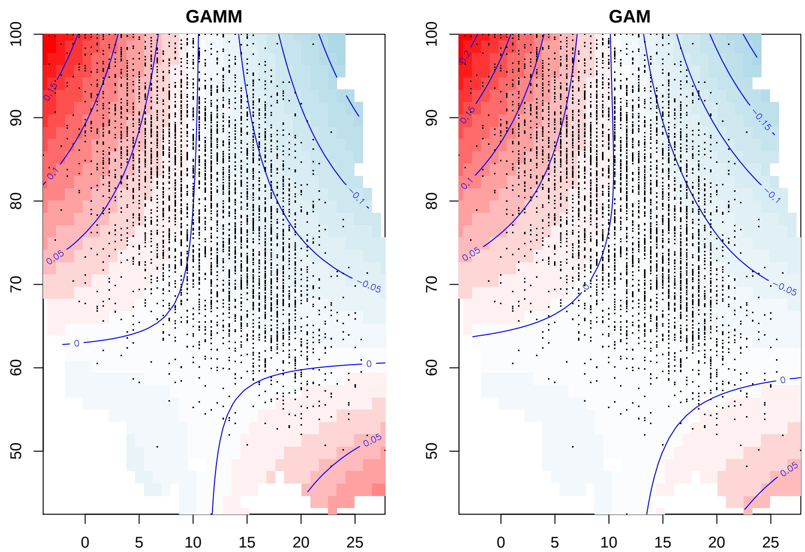

Figure A.10: Bi-variate smooth plots from the fitted generalized additive mixed model with AR(1) residual structure for the lag-4 model (left-hand side), compared to independent errors (right-hand side). As with previous plots, the x-axis is temperature (Celsius) and the y-axes are relative humidity (%).

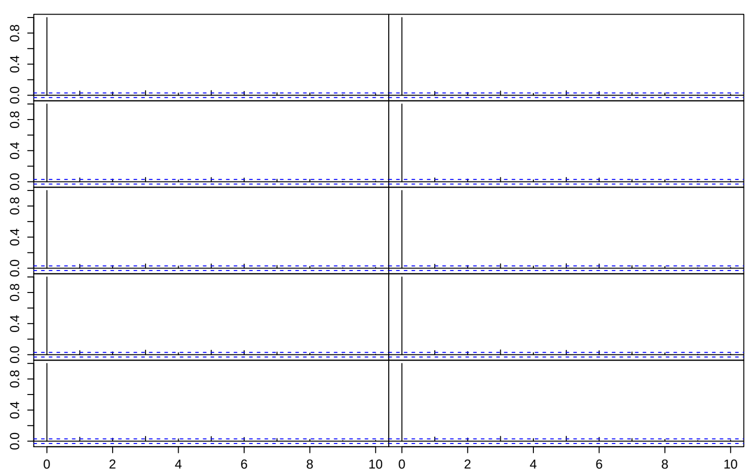

Figure A.11: Residual auto-correlation plots from the lag-k models. On the left-hand side from top to bottom are lags 0, 2, 4, 6 and 8, on the right-hand side are lags 1, 3, 4, 5, 7 and 9.

Table B.2: Parameter estimates for fixed effects.

Model	**Parameter**	**Estimate**	**Standard error**
lag-2	α (Intercept)	1.79	0.01
lag-2	β_1 (Friday)	0.05	0.01
lag-2	β_2 (Monday)	0.01	0.02
lag-2	β_3 (Saturday)	-0.11	0.02
lag-2	β_4 (Sunday)	-0.06	0.02
lag-2	β_5 (Thursday)	0.05	0.01
lag-2	β_6 (Tuesday)	0.04	0.01
lag-2	γ (Public Holiday)	-0.14	0.04
lag-7	α (Intercept)	1.79	0.01
lag-7	β_1 (Friday)	0.05	0.01
lag-7	β_2 (Monday)	0.01	0.02
lag-7	β_3 (Saturday)	-0.11	0.02
lag-7	β_4 (Sunday)	-0.06	0.02
lag-7	β_5 (Thursday)	0.05	0.01
lag-7	β_6 (Tuesday)	0.03	0.01
lag-7	γ (Public Holiday)	-0.16	0.04
Single Index	α (Intercept)	1.79	0.01
Single Index	β_1 (Friday)	0.05	0.02
Single Index	β_2 (Monday)	0.01	0.02
Single Index	β_3 (Saturday)	-0.11	0.02
Single Index	β_4 (Sunday)	-0.06	0.02
Single Index	β_5 (Thursday)	0.05	0.02
Single Index	β_6 (Tuesday)	0.04	0.02
Single Index	γ (Public Holiday)	-0.16	0.05

Effect estimates are shown for the lag-2, lag-7 and single index models. Results for other models were analogous.