A Hierarchical Bayesian Model for Improving Short-Term Forecasting of Hospital Demand by Including Meteorological Information

Sujit K. Sahu[1], Bernard Baffour[2],
Paul R. Harper[3], John H. Minty[3] and Christophe Sarran[4] *

Abstract

The effect of weather on health has been widely researched, and the ability to forecast meteorological events is able to offer valuable insights into the impact on public health services. In addition, better predictions of hospital demand that are more sensitive to fluctuations in weather can allow hospital administrators to optimise resource allocation and service delivery. Using historical hospital admission data and several seasonal and meteorological variables for a site near the hospital, this paper develops a novel Bayesian model for short-term prediction of the numbers of admissions categorised by several factors such as age-group and sex. The proposed model is extended by incorporating the inherent uncertainty in the meteorological forecasts into the predictions for the number of admissions. The methods are illustrated with admissions data obtained from two moderately large hospital trusts in Cardiff and Southampton, in the United Kingdom, each admitting about 30-50 thousand non-elective patients every year. The Bayesian model, computed using Markov chain Monte Carlo methods, is shown to produce more accurate predictions of the number of hospital admissions than those obtained using a six-week moving average method similar to that widely used by the hospital managers. The gains are shown to be substantial during periods of rapid temperature changes, typically during the onset of cold, and highly variable winter weather.

Keywords: Bayesian Inference, Forecast Validation, Emergency Admissions, Prediction.

1 Introduction

There has been a long standing recognition of the relationship between weather and health dating back to the time of Hippocrates, who first discovered that disease was linked to

*1=Southampton Statistical Sciences Research Institute, University of Southampton, Southampton, UK; 2=Institute for Social Science Research, The University of Queensland, St Lucia, Australia. 3=School of Mathematics, Cardiff University, Cardiff, UK, 4=Health Research, Met Office, FitzRoy Road, Exeter, UK. Corresponding Author Email Address: *S.K.Sahu@soton.ac.uk*

changes in weather. In fact, there is a whole scientific discipline, referred to as human biometeorology, which studies the interrelationship between atmospheric conditions and human health, see e.g., Sargeant (1964). In addition, there is a growing wealth of research that suggests that changes in the weather have both direct and indirect influence on human health and/or behaviour. One area of research that has been increasingly investigated is that of excess seasonal mortality and morbidity, in particular during the winter, see, e.g. Curwen (1991), Clinch and Healy (2000), Donaldson and Keatinge (2002), Hajat and Haines (2002), and Healy (2003). Although it is rare for cold weather to kill people directly, cold is indirectly responsible for increased mortality and morbidity through respiratory ailments, notably bronchitis and pneumonia, and through cardiovascular ailments such as coronary heart disease and strokes. This leads to an increase in the demand for hospital services during periods of protracted cold weather. However, various disease conditions have sometimes conflicting relationships with weather. For example, studies have shown that the incidence of coronary heart disease is related to periods of both hot and cold weather (Hajat *et al.*, 2007; Bhaskaran *et al.*, 2010) but the effect of weather on strokes is not clear (Cowperthwaite and Burnett, 2011).

Severe air pollution, in addition to weather, leads to short-term increase in hospital admission rates. In particular, high levels of particulate matter air pollution leads to higher number of admissions of people suffering from chronic obstructive pulmonary disease (COPD), see, e.g. Dominici *et al* (2006) and high levels of atmospheric ozone concentration leads to incresed number of hospital admissions of pediatric asthma patients, see, e.g. Zhu *et al.* (2003). As a result, while developing a forecasting model for the number of hospital admissions we also include air pollution as explanatory variables.

The number of hospital admissions naturally varies with many seasonal variables, such as the day of the week and school holidays. For example, typically more people are admitted on a Monday than a Sunday during normal times (i.e. without an on-going epidemic or major events such as closure of a nearby hospital). These seasonal effects along with many characteristics of weather and their interactions are natural strong contenders for explaining the variability of number of hospital admissions throughout the year. However, these relationships contain many confounders, for example, month and the daily average temperature are highly associated and any analysis using both as explanatory factors will be problematic to fit and interpret.

Accurate forecasting of hospital demand along with their uncertainties offers an indispensable source of information on the demand for public health services. Large National Health Service (NHS) hospitals often have long non-emergency demand for their services, and these will be managed much more efficiently if accurate forecasts for the emergency demand for short-term planning horizons, e.g. up to a week, are available. Even a small improvement in the forecasting method has the ability to save millions of pounds annually because of the huge size of their annual spend on these services. This is likely to reduce the, often, long wait-list for hospital services demanded by non-emergency patients, thereby, improving their experience and long-term survival rates.

Given so much importance and so many potential uses for accurate forecasting, it is only natural to expect the hospitals to have access to a system of robust and well calibrated forecasting method for their short term planning purposes. However, our experience in working with the two large participating hospitals (Southampton and Cardiff) did not find any mathematical or statistical model whatsoever. The operations departments in these hospitals primarily use expert managerial judgments assisted by simple moving averages that are adjusted, often, using the ad-hoc discretion of the managers, for seasonality. Due to the link between influenza outbreak patterns and hospital admissions (see, e.g. Glezen *et al.*, 1982; O'Brien *et al.*, 2000) a six week moving average is preferred to other short and long-term moving averages. The averages are calculated separately for each day of the week and are sometimes found to work reasonably well. However, the six week moving averages (or in fact any other moving average), as forecasts, are not able to take into account variation in number of admissions due to important covariates such as weather. For example, a sudden cold-spell may only lead to a higher number of admissions next week and the moving average cannot predict this sudden surge in demand because it has not yet seen one. Moreover, as in any non-deterministic forecasting method, it is not possible to assess the associated uncertainties of these forecasts. Without these uncertainties to hand, the managers are not able to discover the true extent of the demand.

The main contribution of this paper is the development of an empirical model for short-term forecasting of hospital demand based on the most important meteorological and seasonal variables. The model, required to be simple to implement and interpret, provides accurate forecasts under normal conditions, i.e., without any natural and human catastrophes, throughout the year. Moreover, the model for admissions makes forecasts for number of admissions categorised by sex and age-group of patients, so informing hospitals how better to use resources. For example, current UK government regulations require that adult patients must be admitted in single-sex wards; our model is able to provide forecasts categorised by both sex and age group which will allow the hospitals to plan for their bed capacity much better in line with the requirement. Furthermore, the daily total number of admissions and their uncertainties are automatically obtained as a by-product of the fine level modelling methods proposed here.

Forecasting using the daily demand model that includes past and current weather naturally requires the future values of weather which are not yet observed. An easy solution here is to pretend the forecasts as the observed values. However, this approach will naturally underestimate the uncertainty in the demand forecasts. Our modelling innovation in this paper takes care of this problem by developing a model linking the observed and forecasted meteorological variables. This model, set in a hierarchical Bayesian framework, permits us to propagate and assess the uncertainties in the forecasts for meteorological variables that are being used to forecast the hospital demand. These methods are illustrated and the models are validated using a running window of validation data spanning over the year 2010 from the Southampton General Hospital in the United Kingdom. A smaller data set is also used to confirm the substantive results from a moderately large hospital in Cardiff.

There has been some limited effort in the literature in developing a model for hospital demand. Congdon (2000) develops a geographic model for patient flow by taking care of patient demand, hospital supply and distance effects. Lowthian *et al.* (2012) study 10 year trends in demands at an emergency department. Jones *et al.* (2002) consider forecasting the daily number of occupied beds due to emergency admissions and finds that it is related to air temperature data. However, none of these papers develop a detailed descriptive model for daily number of admissions as is done in this paper.

The outline of the remainder of this paper is as follows. Section 2 provides a brief description of the data. In the modelling Section 3, a large number of explanatory variables are considered to find the best model for hospital demand. Methods for forecasting using the best model are discussed in Section 4. Section 5 develops the unified Bayesian hierarchical model linking demand and meteorological forecasts. Section 6 validates the proposed methods using past hospital admission and meteorological data. Finally, Section 7 provides some concluding remarks.

2 Data description

2.1 Admission Data and Demographic Variables

Individual hospital admission records, anonymised for confidentiality reasons, have been obtained from two participating hospitals in this research project - Southampton University Hospital Trust (SUHT) and Cardiff and Vale University Health Board (CVUHB). SUHT serves a population of around 1.3 million people living in the Southampton and Hampshire area in the south east of England. Owing to its size, it also provides specialist services (such as cardiology and neurosciences) to more than three million people resident in the south of England and the Channel Islands. CVUHB provides health services to a population of around half a million in Cardiff and the Vale of Glamorgan in South Wales. In addition, it provides specialist services (for instance paediatrics care and medical genetics) to the wider population encompassing mid and South Wales, with a population of roughly one million. The analysis is undertaken for both hospitals, but the main results are presented for Southampton and similar results for Cardiff are omitted for brevity.

During the years 2008 and 2009, there have been approximately 94,000 non-elective admissions to the Southampton hospital. Thus there are 127 daily admissions on average that vary considerably by sex and age. It was found that there are more males than females admitted, with there being 65 males and 62 females on average. The average age at admissions is 48 years, in comparison the average age of the UK population is 37 years (Office for National Statistics, 2004). This is consistent with an ageing population since older people are more likely to be hospitalised from age-related conditions such as cancer and cardiovascular diseases.

The hospital operational management does not require the individual ages; instead they work with age split into three groups: *paediatrics* (0-17 year olds), *adults* (18-74 year olds), and *elderly* (aged 75 and above). Therefore, we shall also use these groupings throughout the paper, although it is possible to change the age-groupings according to the needs of individual hospitals. Figure 1 provides the boxplots of the number of admissions for the three age groups further categorised by the day of the week. All three age groups show similar patterns: the number of admissions is much higher during the weekdays than weekend with Mondays receiving the highest numbers of admissions on average.

2.2 Meteorological Data

Daily meteorological data from several sites covering the hospital catchment area have been obtained from the UK's national meteorological service, the Met Office. Ideally, the observed meteorological data where the patients live should only be meaningful for explaining weather related illnesses. However, both the participating hospitals (Southampton and Cardiff) draw in patients from surrounding counties, and even across the country, as they are specialist treatment centres. Moreover, the individual hospital records did not include the postcodes of many patients as many of those were unknown and/or many patients were homeless. That is why we assign the meteorological data for the site covering the actual location of the hospital for all the admissions. This simplification can somewhat degrade the effect of weather variables on hospitalisation, but, bearing in mind that our interest is in developing an aggregated hospital level model for daily number of admissions, this seems to be the best alternative compared to other options such as taking a national average of the meteorological variables. It is also reassuring that more than 80% of the total number of admissions come from a circular area of radius ten kilometres with the hospital at the centre. The meteorological data used in modelling will be correct for these patients.

The effect of temperature on the number of daily admissions categorised by the three age groups is summarised in Figure 2. For all three age groups, on average the number of admissions increases as temperature decreases, although it is worth mentioning that cold weather is not the only factor for hospital admissions as we consider different seasonal variables in the next subsection.

Our analysis also examined the effect of other meteorological variables such as average relative humidity and pressure. However, none of those additional variables explained any further variability in the admission data and are omitted henceforth. We have also examined the effect of severe weather events such as floods, snow and gales but none of those explained any further variability in the data. Lastly, we have considered air pollutants such as ozone concentration levels, particulate matter, sulphur and nitrogen dioxides on the number of hospital admissions in addition to temperature. However, many exploratory analyses conducted to assess their effect on admission did not find any meaningful strong relationships at this *aggregated level* and, as a result, those variables are not considered any further in this paper, although it is clear that these variables will have significant effect on disease specific, such as COPD and asthma, admission rates as mentioned in the Introduction.

2.3 Seasonal Data

Previous analysis makes it clear that the day of the week influences the number of admissions very strongly (Figure 1). It also turns out that the number of admissions is also affected by whether it is a bank holiday. In the UK, there are additional periods when schools are closed, either due to there being a half or end of term. There are two bank holiday Mondays, at the beginning and end of May, and one at the end of August. Also, there are public holidays over Easter (Good Friday and Easter Monday), Christmas (Christmas Day and Boxing Day) and New Year's Day. In addition, there are school closures for a week during the half-terms in February, June and October. Schools are also closed for approximately two weeks around Easter, five weeks during July and August and another two weeks during the Christmas and New Year holiday period.

We investigated the effect of holidays on admissions, and found that although there were slight peaks in admissions during Bank holiday weekends, using school holidays showed the effect more clearly. The number of daily admissions drops considerably for the paediatrics and adults but not for the elderly patients during the school holiday periods, see Figure 3. This could possibly be explained by the fact that during school holidays, families with school-going children spend vacations often away from their homes.

The numbers of admissions of the paediatrics and elderly age groups, unlike those in the adult age group, are higher during the winter months of November to March compared to those during the other months (see Figure 4). As is expected, Figure 4 also reveals that the average temperature is lower during the winter months than the summer months. That leads to strong collinearity between the month, treated as a factor, and the continuous explanatory variable temperature in our model for the number of admissions. However, of primary interest here is to forecast dis-aggregated daily number of admissions and not overall monthly totals. Hence, we do not include month as one of the explanatory factors and build our models with daily temperature instead. Seasonal information, such as public and school holidays, contained in the variable month, already enters our model explicitly through the seasonal variables described in this section.

3 Model Building and Selection

The response variable to model is the number of daily admissions classified by three age categories and sex. Our main aim here is to build a good explanatory and a predictive model based on the available explanatory variables, i.e. the demographic, meteorological and the seasonal variables. The total number of such variables, their interactions and their lagged effects is very large and as a result model building, fitting and selection using multivariate time series models are computationally prohibitive and are not pursued here. Although it is certainly possible to experiment with multivariate and independent univariate time series models with a small number of seemingly important explanatory variables such as temperature. Instead, we adopt multiple regression models and use model selection criteria such as the Akaike Information Criteria, AIC (Akaike, 1974) and the Bayesian Information Criteria, (BIC) (Schwartz, 1978) to select the best model among all the competing models. In passing, we note that an alternative to the model selection approach here will be Bayesian model averaging for forecasting (Raftery *et al.*, 2005), which may be investigated in futute.

Often, the first step in regression modelling is selecting the scale of the response variable. Here the response variable, the daily number of hospital admissions, categorised by age and sex, is a discrete count variable and can theoretically take the zero value – a reason why we avoid the log transformation. The square-root transformation stabilises the variance for count data for which the Poisson distribution is appropriate, see e.g. Box and Cox (1964). Hence, we adopt this transformation that also encourages symmetry and normality of the residuals. Indeed, the residuals under the square-root transformation were much better behaved than those when data were modelled on the original scale. However, we report all the predictions and their uncertainties on the original scale for ease of use and interpretation by the hospital managers. This allows us to avoid implementing a Poisson generalised linear model that only complicates model fitting but does not provide improved fit and forecasting as evidenced by our analyses, which have been omitted for brevity.

The models are developed using the preceding one year of daily data to capture the effects of seasonality in temperature, school holidays and other variables. Of course, accuracy can be increased by modelling data further in the past. However, one of our main task here is to minimise the amount of data needed to train the models without sacrificing forecast accuracy. Modelling using one year of past data was found to be the best option under these considerations. In Section 6 we provide evidence of accurate forecasting based on several data sets each for a running window of one year.

We define the model as follows. Let Y_{ijt} denote the number of admissions on the square root scale for the ith sex ($i = 1$ for female and 2 for male), jth age group ($j = 1, 2, 3$ respectively for paediatrics, adult and elderly) and tth day where $t = 1, \ldots, 365 (= T)$ corresponding to one year's data for modelling. We also introduce the matrices $\mathbf{X}_{1t}$, $\mathbf{X}_{2t}$ and $\mathbf{X}_{3t}$, each of appropriate order, that collect the demographic, meteorological and seasonal variables, respectively. In our modelling $\mathbf{X}_{1t}$ contains the age and sex information. The meteorological variables, $\mathbf{X}_{2t}$, contain a large number of variables pertaining to the weather (such as temperature, humidity, rainfall and wind speed) and air quality (such as the levels of carbon monoxide, ozone and particulate matter in the atmosphere). Since there is often a lag between weather conditions and mortality/morbidity we also looked at the lagged daily values up to three weeks. The seasonal data, $\mathbf{X}_{3t}$, consist of temporal information including the day of the week and month, and other seasonal information such as the season and whether or not it was a school holiday or bank holiday.

The most general form of our regression model is:

$$Y_{ijt} = g(\mathbf{X}_{1t}, \mathbf{X}_{2t}, \mathbf{X}_{3t}) + \epsilon_{ijt}, \tag{1}$$

for $i = 1, 2$ and $j = 1, 2, 3$ and $t = 1, \ldots, T$, where $g(\cdot)$ is a general regression function that may include all the variables, $\mathbf{X}_{1t}, \mathbf{X}_{2t}, \mathbf{X}_{3t}$ and their first and second order interactions, and ϵ_{ijt} is the error term. The general model (1) constitutes a huge number of possible regression models including all the demographic, meteorological, and seasonal variables. We assume that these variables are able to take care of the temporal dependence in the daily data and hence we assume the errors, ϵ_{ijt} to be independently and identically distributed as Gaussian random variables with mean 0 and unknown variance σ^2. This independence assumption will be diagnosed by performing appropriate residual analyses, see below.

These multiple regression models, being non-Bayesian, facilitate faster computation so that the initial screening can be performed within a reasonable amount of computing effort and time. Bayesian modelling of the data, especially for inference and forecasting purposes, proceeds with the chosen explanatory variables as is a common practice in the literature, see e.g. Sahu *et al.* (2007) for an ozone concentration modelling problem with many meteorological variables as the explanatory variables.

An extensive model search was performed using many step-wise model selection procedures using the BIC criterion for the initial screening of the explanatory variables. In all investigations, the age-group turns out to be the most important explanatory factor, followed by sex, day of week, the minimum temperature last week and mean daily temperature. In addition, the model with the interaction terms age-sex and age-minimum temperature last week was found to be the best according to the BIC criterion. The interaction terms improved the descriptive capability of the model significantly and the multiple R^2 increased from 76.50% for the main effects only model to 80% for the interaction model fitted to the Southampton data. For the Cardiff data set the multiple R^2 improved from 85.35% to 90% showing a very good model fit.

As in any regression modelling, an investigation was taken into the behaviour of the residuals using diagnostics plots. Two most important diagnostic plots, viz. the Anscombe residual plot and the normal probability plot did not show any unusual patterns that could suggest any substantial departure from the model assumptions. This was to be expected since the data had already been transformed to stabilise the variance and to encourage symmetry.

A further important issue in the residual analysis here is to check for serial correlation since the modelled data are temporally correlated as mentioned above. Toward this end we have obtained the lagged residual plots (shown in Figure 5), plotting r_{ijt} against r_{ijt-1} for each of six combinations of i and j corresponding to two levels of sex and three levels of age where r_{ijt} denotes the residual corresponding to ϵ_{ijt}. None of these plots showed any trend or pattern and all of them showed a random scatter. Thus the independence assumption in the model is deemed to be justified and henceforth we proceed with this model.

The best fitted regression model is given by

$$Y_{ijt} \sim N(\mu_{ijt}, \sigma^2), \tag{2}$$

where

$$\mu_{ijt} = \beta_0 + \alpha_i + \gamma_j + h(t) + w(t) + \lambda_m m_t + \lambda_n n_t + (\alpha : \gamma)_{ij} + (\gamma : n_t)_j.$$

Here β_0 is the intercept term; α_i is the additional intercept for sex and it is zero for females; γ_j is the effect of the jth age group and it is zero for paediatrics; $h(t)$ is the school holiday effect and it is zero for non-school holidays, $w(t)$ is the day of the week effect and it is zero for Sundays; λ_m is the coefficient of mean temperature m_t on day t; λ_n is the coefficient of minimum temperature n_t on day $t-7$ i.e. a week before; $(\alpha : \gamma)_{ij}$ is the age-sex interaction term; and $(\gamma : n_t)_j$ is the age and minimum temperature last week interaction term. In the two interaction terms, the parameter value is constrained to be zero when any of the factors is at their first level, e.g. $(\gamma : n_t)_1 = 0$. Thus, the model is described by 18 regression parameters β_0, α_2, γ_2, γ_3, one parameter $h(t)$ when the day t is not a school holiday, six $w(t)$ for the six days of the week (Monday-Saturday) $w(t) = 0$ when t corresponds to a Sunday;

$(\alpha : \gamma)_{22}$,$(\alpha : \gamma)_{23}$, $(\gamma : n_t)_2$ and $(\gamma : n_t)_3$. Collectively we denote these 18 parameters by $\boldsymbol{\beta}$ and denote the regression model (2) by the notation

$$\mathbf{Y} = X\boldsymbol{\beta} + \boldsymbol{\epsilon}$$

by stacking the individual Y_{ijt}'s and the ϵ_{ijt}'s in any suitable order. Here X denotes the resulting design matrix.

Table 1 provides the parameter estimates of the best model fitted separately to the Southampton and Cardiff data. The main effects, due to sex, day of the week, and holiday, are broadly similar for the models fitted to data from two hospitals, pointing to perhaps a broad socio-economic phenomenon. However, the Cardiff hospital admits a significantly smaller proportion of elderly than paediatric patients. Daily mean temperature and minimum temperature last week are both significant for the Southampton data, but they are not significant for the Cardiff data. Nevertheless, the interaction effect of age group and temperature last week is significant for both the data sets. Hence, we keep this effect and the corresponding main effects in the model. We keep the daily mean temperature in the model for Cardiff data for the sake of convenience in having to fit and forecast one model for both the hospitals.

4 Forecasting

Recall that our main objective is to provide forecasts of number of admissions up to seven days in advance. According to the best model (2), the forecast distribution of the number of admissions, looking ahead at day T, for k days in advance, $k = 1, \ldots, 7$ has the mean given by

$$\mu_{ijT+k} = \beta_0 + \alpha_i + \gamma_j + h(T+k) + w(T+k) + \lambda_m m_{T+k} + \lambda_n n_{T+k} + (\alpha : \gamma)_{ij} + (\gamma : n_{T+k})_j. \quad (3)$$

The plug-in approach, under classical inference methods, simply replaces the parameters by their estimates in the above forecasting mean and reports the forecast by the corresponding estimate $\hat{\mu}_{ijT+k}$. The forecast standard error can also be calculated using standard methods.

There is a problem, however, when forecasting in a typical real time situation, exactly where we envision the primary use of these models. Calculation of the forecasts will require us to know the observed mean temperatures for the *future* days, $T+k$, for $k = 1, \ldots, 7$. Obviously, this is impossible. Note that this problem does not arise for the minimum temperature last week, n_{T+k}, since n_{T+k} is the minimum temperature on day $T + k - 7$ which has been observed already, for all $k = 1, \ldots, 7$ at day T. Returning to the problem, a straightforward solution is to replace the m_{T+k} by the k-day ahead forecasted temperature, $f_T^{(k)}$ say, provided by the Met Office on current day T. However, as can be expected there will be differences between the observed and the forecasted temperatures, and, usually, uncertainty in the forecasted temperatures $f_T^{(k)}$ will increase for increasing values of k. These uncertainties in the forecasted temperatures must be propagated into the forecasts for the number of admissions using sound methodologies.

The estimation of the uncertainties in the forecasted temperatures and their refinements is possible with the help of past observed data for a number of days and the forecasts that were made for them. Suppose that we have the observed mean temperatures, m_t and its k-day ahead forecast $f_{t-k}^{(k)}$ made at day $t-k$ for each $k = 1, \ldots, 7$ and $t = 1, \ldots, L$ days. The dates for these data do not have to coincide with the dates for the admission data, and as a result L does not have to be equal to T. This is an auxiliary data set that will be used to refine the forecasts and to propagate the uncertainty in these forecasts into the admissions model. In fact, in our hierarchical modelling we will assume that these data are independent of the admissions data and the meteorological data used in the model (2).

The forecasts of the temperature, being output of numerical computer simulation model, are sometimes biased and a simple bias correction model is given by:

$$m_t = a_k + b_k f_{t-k}^{(k)} + \eta_t^{(k)}, \quad t = 1, \ldots, L, \tag{4}$$

for each of $k = 1, \ldots, 7$, where we assume that the errors $\eta_t^{(k)} \sim N(0, \tau_k^2)$ independently. The parameters a_k and b_k measure the additive and the multiplicative biases in the forecast $f_{t-k}^{(k)}$ for the observation m_t, see e.g. Fuentes and Raftery (2005) regarding the biases in a computer model. The parameter τ_k^2 corresponds to the uncertainty in the residuals and may increase with k.

Our proposal here is to fit model (4) for each k separately, possibly with available independent past data, and thereby to estimate a_k and b_k along with their uncertainties. Once this modelling has been performed, we can replace the unobserved m_{T+k} by the refined forecasts

$$\hat{m}_{T+k} = \hat{a}_k + \hat{b}_k f_T^{(k)}, \quad k = 1, \ldots, 7 \tag{5}$$

in the admission forecasting model (3), where $\hat{a}_k$ and $\hat{b}_k$ denote the estimates of a_k and b_k. This method, however, will not incorporate the uncertainties in the parameter estimates $\hat{a}_k$ and $\hat{b}_k$ into the forecasts for the number of admissions. In fact, it is not straightforward to propagate these uncertainties in a coherent manner using classical inference methods. Instead, we propose to perform this task using a Bayesian hierarchical model in the next section.

5 Hierarchical Bayesian Modelling

5.1 Model Specification

At the top level of the Bayesian hierarchy, we continue to assume the admission model (2), and in the following discussion use the simplified notation $\mathbf{Y} \sim N(X\boldsymbol{\beta}, \sigma^2 I)$ where I is the identity matrix of appropriate order. This Bayesian model is completed by assuming suitable prior distributions for $\boldsymbol{\beta}$ and σ^2. For $\boldsymbol{\beta}$ we assume the default vague prior where each component is assumed to be independently normally distributed with mean 0 and a large variance 10^4. We assume that the precision parameter $1/\sigma^2$ follows the gamma distribution with parameters $\nu > 0$ and $\xi > 0$ where the gamma distribution has mean ν/ξ. In our implementation, we choose $\nu = 2$ and $\xi = 1$ to have a proper prior specification that avoids the pitfalls of improper prior distributions, see e.g. Sahu *et al.* (2007). This prior distribution will be assumed for all variance parameters in the modelling development below.

Independently of the admissions model we assume model (4) for the observed temperatures based on the forecasts. As in the above paragraph, we assign independent normal prior distribution with mean 0 and variance 10^4 for each of the parameters a_k and b_k, for $k = 1, \ldots, 7$. The parameters $1/\tau_k^2$ are given independent gamma prior distributions each with hyper-parameter ν and ξ.

The joint posterior distribution of all the parameters $\boldsymbol{\beta}, \sigma^2, a_k, b_k$, and τ_k^2 for $k = 1, \ldots, 7$, denoted by $\boldsymbol{\theta}$ is obtained from the above specifications using standard methods, the details are omitted. It is important to note that there is no possibility of feedback from the temperature observation and forecast model (4) to the admission model (2) since the posterior distributions of the parameters under each model are calculated independently.

Table 2 provides the Bayes estimates of a_k, b_k and τ_k^2, for $k = 1, \ldots, 7$ fitted to the observed temperature and forecasts data for Southampton for the 365 days in the year 2010. The estimates show that the 1-3 days ahead forecasted temperatures are significantly upwardly biased, perhaps due to the small data set from a single site that has been used for this model fitting. The bias is not significant for the 4-7 days ahead forecasts. Moreover, as expected there is no significant multiplicative bias since the 95% credible intervals for b_k all include the value 1. The estimates of the variances, $\tau_k^2, k = 1, \ldots, 7$ are expected to increase with k, but, on the contrary, a dip is observed for the 3-day ahead forecasts. This is explained as follows. The forecasts for days 1 and 2 come from the Met Office's Global Unified Model and are deterministic forecasts. The forecasts for days 4 and beyond come from the European Community's model ensemble mean. The forecast for day 3 is a mixture (blend) of the two sources of forecasts. It has been found that the blending of models provide much better forecasts, see e.g. Evans *et al.* (2000) and Bowler *et al.* (2008).

5.2 Forecasting

The adopted Bayesian paradigm is particularly attractive for the task of forecasting the number of admissions. Here each forecast will have a posterior predictive distribution that is derived as follows. First, note that the forecast distribution for Y_{ijT+k}, conditional on the parameters, will be $N(\mu_{ijT+k}, \sigma^2)$ where μ_{ijT+k} is given by (3). The mean μ_{ijT+k} depends on the future temperature observation m_{T+k} which is specified, according to (4), to be:

$$m_{T+k} \sim N(a_k + b_k f_T^{(k)}, \tau_k^2), \quad k = 1, \ldots, 7 \tag{6}$$

conditional on the parameters a_k, b_k and τ_k^2.

Second, the posterior estimates of a_k, b_k and τ_k^2 along with the Met Office provided $f_T^{(k)}$ will enable computation of the posterior predictive distribution of m_{T+k} given the past temperature and forecast data. This posterior predictive distribution of m_{T+k} will then need to be combined with the posterior distribution of $\boldsymbol{\beta}$ and σ^2 to obtain the posterior predictive distribution of Y_{ijT+k} given all the past admission, meteorological and seasonal data.

A quick plug-in approach, that has an obvious analogue in classical inference methods, is to use the posterior means of a_k and b_k, say $\hat{a}_k$ and $\hat{b}_k$ respectively, in (5), and then proceed with the admission forecasting model (3) replacing m_{T+k} by $\hat{m}_{T+k}$. This method, however, will not be able to assess the uncertainties in the forecasts and that is why we adopt the following Monte Carlo simulation method.

The posterior predictive distribution of the forecast Y_{ijT+k} is obtained by composition sampling based on Markov Chain Monte Carlo (MCMC) methods as follows. We first obtain the samples $\boldsymbol{\theta}^{(s)}$, $s = 1, \ldots, S$ for a large value of S from the joint posterior distribution of $\boldsymbol{\theta}$ given all the data. Then for each s we obtain

$$m_{T+k}^{(s)} \sim N\left(a_k^{(s)} + b_k^{(s)} f_t^{(k)}, \tau_k^{2(s)}\right), k = 1, \ldots, 7.$$

At the next step we obtain

$$Y_{ijT+k}^{(s)} \sim N\left(\mu_{ijT+k}^{(s)}, \sigma^{2(s)}\right), k = 1, \ldots, 7$$

where $\mu_{ijT+k}^{(s)}$ is obtained from (3) after replacing the parameters by their simulated values in $\boldsymbol{\theta}^{(s)}$, and m_{T+k} by $m_{T+k}^{(s)}$. Once, we have the samples $Y_{ijT+k}^{(s)}$, $s = 1, \ldots, S$ we transform these back to the original scale (here by simply squaring each value). The MCMC samples on the original scale are averaged to estimate the forecast mean of Y_{ijT+k}. The samples are also used appropriately to obtain the standard errors of the averages and the 95% forecast intervals. In our implementation we take $S = 5000$ after discarding the first 1000 iterates, although MCMC convergence is not an issue at all. Moreover, the computations are fast because of conjugacy.

6 Analysis and Forecast Validation

The success of the proposed methodology is to be judged wholly by the performance of the validation forecasts for hold-out data spread over an entire year – so that the model can be trusted for a whole calendar year that includes all the seasons. However, it is also clear that the considered year may miss severe weather, by chance, and additional experimentation may be necessary. In fact, this is recommended in the discussion Section 7 where it is also emphasised that the model be re-calibrated with most recent data as much as possible.

Our validation method with one year's data requires us to use a moving window of one year's data for modelling and then comparing the forecasted number of admissions with the observed ones, categorised by the six age-group and sex combinations, for the next seven days. At each day, the previous year's data are used to fit the model, and then we predict ahead for the next seven days. Thus, when the moving window advances by just a day, ideally, we need to compare the 15,330 ($= 365 \times 6 \times 7$) forecasts with the 2,190 ($= 365 \times 6$) observed values.

However, for the Southampton hospital we only have the full meteorological data, i.e. both the observed and forecasted temperatures, for the year 2010. Hence, the last day we stop data collection is 24th of December, so that we have all the comparison data available for seven days ahead until December 31. As a result, our validation comparisons are based on actual data from 359 days.

For the Cardiff data set, we do not have the forecasted meteorological data but have the observed temperature data and the admissions data from 25th April 2010 to January 31, 2012. Hence, for forecast validation purposes, the first forecasting window starts on April 25, 2011 (allowing for previous one year's data for modelling), and as a result we compare forecasts for 278 days (from April 25, 2011 to January 31, 2012) with the corresponding observed ones.

For both hospitals, we compare the following three forecasting methods. The first one is the persistence forecast, which is simply the most recent observation, and in our application refers to the number of admissions a week ago. The persistence forecasts, simplest to obtain, present a baseline method of forecasting and facilitate comparison. The second method uses a six-week moving average method, similar to that currently used by the hospital operations departments. The third forecasting method treats the observed temperatures as the forecasts temperature, and thus ignores the inherent uncertainty associated with the forecasts.

We compare the above three methods with the the proposed full Bayesian method that takes care of the uncertainties in the forecasts. However, the temperature forecasts for the Cardiff hospital are not available for the entire period, and hence we do not consider the last method for the Cardiff data set.

In this paper, to compare the above forecasting methods, we use four key summary criteria: the root mean square error (RMSE), the mean absolute error (MAE), the relative bias (RB), and the nominal coverage. A good forecasting method should have low RMSE and MAE values, which are on the original unit of the data. The RB should be close to zero for forecast unbiasedness. Lastly, the achieved coverage should be close to its true value. This last criterion is very stringent as it not only requires the central tendency of the Bayesian forecast distribution to be correct, but also demands their correct uncertainty assessment, see e.g. Gneiting *et al.* (2007).

We now consider validation of the model using data for the whole year, for both the hospitals. We compare the predictive performance of the forecasted number of admissions for our proposed method against the six week moving average and the persistence. Tables 3 and 4 provide the RMSE, MAE and RB for all three methods of forecasting. In addition, the nominal coverages of the 95% forecast intervals for the proposed Bayesian model are shown; such coverages are not meaningful for the persistence forecasts and six-week moving averages. As expected, the persistence method fares worst in terms of the RMSE, MAE and RB. This is not surprising since it is a fairly naive forecasting approach and fails to account for any trends in admissions, unlike the moving average and the proposed Bayesian model. Further, the first three criteria show that, overall, the Bayesian methods are uniformly better than the six-week moving average method. In particular, the relative biases of the proposed methods are almost half of that for the six-week moving averages. Typically, there is almost 3-6%

reduction in the RMSE and MAE as well. This reduction in RMSE and MAE is similar for both Southampton and Cardiff hospitals.

Note that the RMSE for the Bayesian method when observed temperatures are plugged-in does not increase with the forecasting step (number of days). This is expected since the model (2) is not a time-series model, as mentioned in Section 3. However, as expected both the RMSE and MAE increase when the forecasted temperatures are used instead of the observed temperature. The RMSEs and the MAEs for the forecasted temperature method are close to those for the observed temperature method which shows that the forecasts and the observed temperatures are also close.

Figure 6 provides a plot of the RMSE for each day during the entire validation period for each of the two hospitals. As expected, both the six-week moving average forecasts and the model based forecasts, using observed temperature data, outperform the persistence method. According to the RMSE, the model-based method outperforms the six-week moving average method on the vast majority of days. Remarkably, the top panel of this figure, for the Southampton data, also reveals that a dip in the temperature graph is followed by a dip in the RMSE for the model based method and the dipped RMSE values are smaller than the same for the six week moving average method. The same phenomenon is also seen in the plot for Cardiff data (bottom panel) except for the first half of the month of November, which is perhaps due to the effect of some unusual events, e.g. closure of a nearby hospital, during the time at Cardiff. This analysis demonstrates that the proposed model based method is likely to be more accurate during rapid temperature changes, typically during the onset of cold, and highly variable winter weather.

The nominal coverages, both overall and the ones broken down by the number of forecasting days ahead, of both the methods are very near to the true value of 95%. This shows that the forecasts using the proposed methods have the correct amount of uncertainties associated with them. It is also remarkable that these measures are based on validation data spanning over a year, hence the proposed model and the forecasting methods are likely to hold true throughout the year.

The lengths of the 50% and 90% forecast intervals are shown as boxplots in Figure 7 for both of the proposed methods. As expected, the forecasted temperature method has slightly more variability in length than the observed temperature method, thus some intervals under the forecasted temperature method will be wider than the ones under the observed temperature method.

The above overall analyses are followed up by more detailed analysis where the forecasts are broken down by the day of the week and whether it is a school holiday or not, see Figures 8 and 9. As in the previous overall case, the two figures show that the RMSE and MAE are higher for the forecasted temperature method than the observed temperature method. More importantly, as can be expected, the models find it easier to forecast for Saturday and Sunday than for the weekdays. However, there seems to be no large differences in all four criteria values for the weekdays and also for weekends. There is no pattern either in the nominal coverages or the relative biases for the seven days further classified by school holiday. These results show the overall stability of the forecasts.

7 Discussion

This paper shows that forecasting of number of hospital admissions is significantly improved by incorporating meteorological information. A parsimonious model for number of daily admissions is obtained by including the mean temperature of the day and the minimum observed temperature last week at a site near the hospital. The model is thus able to adjust to both immediate and past changes in temperature, and has been shown to perform particularly well when there are sudden changes in temperature which may occur, for example, during the onset of cold weather during the autumn. The model also includes two most important demographic variables: age-group of patients and their gender and thus is able to produce forecasts of number of admissions for each of the sex and age group categories. The two most important seasonal variables, day of the week and whether it is a school holiday, are also seen to be significant in the final model that includes interaction between minimum temperature last week and age-group. This is intuitively justified by the fact that the minimum temperature has a different effect on different age groups.

A unified Bayesian hierarchical model has been developed for the short-term forecasting purposes. The hierarchical feature of the model enables us to propagate the uncertainties in the meteorological forecasts into the overall uncertainties in the forecasts for number of admissions. An extensive investigation using several forecast validation criteria shows the superiority of the methods over the simple moving average method of forecasting often used by the hospital managers. These results have been observed for admissions data for both the participating hospitals.

The Bayesian model can be extended to provide estimates of daily admissions by specific ailments as well. This has been investigated, but results are not reported because of the presence of a large number of zero counts in the data finely categorised by age-group, sex and ailments. A possible extension is to pool the data from different hospitals thereby avoiding the overwhelming number of zero counts. A hospital specific random effect can be included to account for differences between hospitals. Another possibility is to model aggregated data at a lower level temporal resolution, e.g. the weekly counts instead of the dailies.

Returning to the cost-saving point made in the Introduction, it is anticipated that the improved forecasts will enable hospitals to react to predicted changes in emergency demand and then to adjust their elective schedules more efficiently. Indeed, anecdotal evidence suggests this to be the case for the two participating hospitals. We expect to more formally capture cost-benefits savings through further planned pilots with an increased number of hospitals over the winter 2012/13.

The proposed method is only designed to work under normal operating circumstances. It will not be able to cope with catastrophes caused either by nature or human. The model is also not likely to work unchecked and automatically for long periods of time. It is recommended that the model is re-calibrated with most-recent data as frequently as possible, e.g. monthly, so that recent changes in admission practices and any systematic errors are dealt with in a regular time interval. The auxiliary observed temperature and forecast data are also need to be updated so that any improvement in weather forecasting is fed to the demand forecast model. Although the model has been shown to produce good forecasting results for the two participating hospitals, it is recommended that a fresh model search, as done in Section 3, is performed for each hospital so that any local effect is taken into account in the best possible way.

Acknowledgement

The authors thank Patrick Sachon from the Met Office; Alison Gray and Ben Hendy from the Southampton University Hospital Trust; and Andrew Nelson, Steve Tarr and John Peters from the Cardiff and Vale University Health Board for providing the data used here and sharing their many helpful insights.

References

Akaike, H. (1974). A new look at the statistical model identification. *IEEE Transactions on Automatic Control*, **19**(6), 716–723.

Bhaskaran, K., Hajat, S., Haines, A., Herrett, E., Wilkinson, P. and Smeeth, L. (2010). Short term effects of temperature on the risk of myocardial infarction in England and Wales: times series regression analysis of the myocardial ischaemia national audit project (MINAP) registry, *British Medical Journal (BMJ)*, **341**(3823), 1136-1146.

Bowler, N. E., Arribas, A. and Mylne, K. R (2008). The Benefits of Multianalysis and Poor Man's Ensembles. *Monthly Weather Review*, **136**, 4113-4129.

Box, G.E.P. and Cox, D.R. (1964). An analysis of transformations (with discussion). *Journal of the Royal Statistical Society, Series B*, **26**, 211–252.

Clinch, J.P. and Healy, J.D. (2000). Housing standards and excess winter mortality. *Journal of Epidemiological Community Health*, **54**, 719–720.

Congdon, P. (2000) A Bayesian approach to prediction using the gravity model, with an application to patient flow modeling. *Geographical Analysis*, **32**, 205-224.

Cowperthwaite, M.C. and Burnett, M.G. (2011). The association between weather and spontaneous subarachnoid hemorrhage: An analysis of 155 US hospitals. *Neurosurgery*, **68**(1), 132–139.

Curwen, M. (1991). Excess Winter Mortality: a British phenomenon? *Health Trends*, **22**, 169–175.

Dominici, F., Peng, R. D., Bell, M. L., Pham, L., McDermott, A., Zeger, S. L. and Samet, J. M. (2006). Fine particulate air pollution and hospital admission for cardiovascular and respiratory diseases. *Journal of the American Medical Association*, **295**, 1127-1134.

Donaldson, G. and Keatinge, W. (2002). Excess Winter Mortality: Influenza or Cold stress? Observational Study. *British Medical Journal, (BMJ)*, **324**(7329), 89–90.

Evans, R. E., Harrison, M. S. J., Graham, R. J. and Mylne, K. R. (2000). Joint medium-range ensembles from the Meteorological Office and ECMWF systems. *Monthly Weather Review*, **128**, 3104-3127.

Fuentes, M. and Raftery, A. (2005). Model Evaluation and Spatial Interpolation by Bayesian Combination of Observations with outputs from Numerical Models. *Biometrics*, **61**, 36–45.

Glezen, P.W., Payne, A.A., Snyder, D.N. and Downs, T.D. (1982). Mortality and Influenza. *Journal of Infectious Diseases*, **146**(3), 313–321.

Gneiting, T., Balabdaoui, F., and Raftery, A.E. (2007). Probabilistic forecasts, calibration and sharpness. *Journal of the Royal Statistical Society. Series B: Statistical Methodology*, **69**(2), 242–268.

Hajat, S. and Haines, A. (2002). Associations of cold temperatures with GP consultations for respiratory and cardiovascular disease amongst the elderly in London. *International Journal of Epidemiology*, **31**(4), 825–830.

Hajat S., Kovats R.S. and Lachowycz K. (2007). Heat-related and cold-related deaths in England and Wales: who is at risk? *Occupational and Environmental Medicine*, **64**(2), 93–100.

Healy, J. (2003). Excess winter mortality in Europe: a cross country analysis identifying key risk factors. *Journal of Epidemiology and Community Health*, **57**, 784–789.

Jones, S. A., Joy, M. P. and Pearson, J. (2002) Forecasting demand of emergency care. *Health care management science*, **5**, 297-305.

Lowthian, J. A., Curtis, A. J., Jolley, D. J., Stoelwinder, J.U., McNeil, J. J., and Cameron, P. A. (2012) Demand at the emergency department front door: 10-year trends in presentations. *Medical Journal of Australia*, **196**, 128-132.

Office for National Statistics (2004). National Statistics Online - Population Estimates. Available at www.neighbourhood.statistics.gov.uk/dissemination. Accessed on 01/04/2012.

O'Brien, K.L., Walters, M.I., Sellman, J, Quinlisk, P., Regnery, H., Schwartz, B. and Dowell, S.F, (2000). Severe Pneumococcal Pneumonia in Previously Healthy Children: The Role of Preceding Influenza Infection. *Clinical Infectious Diseases*, **30**(5), 784–789.

Raftery, A. E., Gneiting, T., Balabdaoui, F. and Polakowsk, M. (2005). Using Bayesian Model Averaging to Calibrate Forecast Ensembles. *Monthly Weather Review*, **138**, 4199-4211.

Sahu, S. K., Gelfand, A. E. and Holland, D. M. (2007). High Resolution Space-Time Ozone Modeling for Assessing Trends. *Journal of the American Statistical Association*, **102**, 1221–1234.

Table 1: Parameter estimates for the best fitting model for the Southampton and Cardiff hospital data. The model R-squared values are 0.806 for Southampton and 0.910 for Cardiff, with residual standard errors of 0 .5313 and 0.5477, respectively. For Southampton, the AIC=3464.11 and BIC=3566.56; while for Cardiff, the AIC=3603.77 and BIC=3706.22.

	Southampton		**Cardiff**	
Parameter	Estimate	95% CI	Estimate	95% CI
Intercept	3.293	(3.189, 3.398)	4.389	(4.196, 4.439)
Male (α_2)	0.576	(0.499, 0.653)	0.359	(0.279, 0.439)
Adult (γ_2)	2.380	(2.260, 2.501)	3.169	(3.175, 3.435)
Elderly (γ_3)	0.903	(0.782, 1.023)	–1.676	(–1.726, –1.466)
Monday	0.446	(0.362, 0.529)	0.664	(0.578, 0.751)
Tuesday	0.316	(0.233, 0.400)	0.740	(0.653, 0.836)
Wednesday	0.377	(0.293, 0.460)	0.703	(0.617, 0.789)
Thursday	0.302	(0.219, 0.385)	0.763	(0.677, 0.849)
Friday	0.348	(0.265, 0.432)	0.806	(0.719, 0.892)
Saturday	–0.042	(–0.125, 0.042)	0.170	(0.084, 0.256)
Holiday	–0.012	(–0.169, –0.066)	–0.095	(–0.148, –0.042)
Daily Temperature (λ_m)	–0.011	(–0.018, –0.005)	0.014	(–0.033, 0.000)
Temperature Last Week (λ_n)	–0.015	(–0.024, –0.006)	–0.034	(–0.022, 0.016)
Male:Adult ($(\alpha : \gamma)_{22}$)	–0.331	(–0.440, –0.222)	–3.749	(–3.861, –3.636)
Male:Elderly ($(\alpha : \gamma)_{23}$)	–1.149	(–1.258, –1.040)	–0.680	(–0.793, –0.568)
Adult:Temp. Last Week ($(\gamma : n_t)_2$)	0.015	(0.004, 0.026)	0.032	(0.014, 0.041)
Elderly:Temp. Last Week ($(\gamma : n_t)_3$)	0.004	(–0.007, 0.015)	0.023	(0.008, 0.035)

Sahu, S. K., Yip, S., and Holland, D. (2009). Improved space-time forecasting of next day ozone concentrations in the eastern U.S. *Atmospheric Environment*, **43**(3), 494–501.

Schwarz, G. E. (1978). Estimating the dimension of a model. *Annals of Statistics*, **6**, 461–464.

Sargeant, F. (1964). A survey of human biometeorology. *International Journal of Biometeorology*, **9**(1), 1–3.

Zhu, L. Carlin, B. P. and Gelfand, A. E. (2003). Hierarchical regression with misaligned spatial data: relating ambient ozone and pediatric asthma ER visits in Atlanta. *Environmetrics*, **14**, 537-557.

Table 2: The Bayes estimates of parameters in the temperature forecast model, a_k, b_k and τ_k^2, for $k = 1, \ldots, 7$, with 95% credible intervals.

day	a_k		b_k		τ_k^2	
	Estimate	95% CI	Estimate	95% CI	Estimate	95% CI
1	–1.19	(–1.48, –0.91)	1.01	(0.98, 1.04)	2.27	(1.95, 2.65)
2	–1.16	(–1.46, –0.86)	1.01	(0.98, 1.04)	2.41	(2.07, 2.81)
3	–0.57	(–0.82, –0.33)	1.03	(1.00, 1.05)	1.79	(1.53, 2.09)
4	0.11	(–0.16, 0.57)	1.02	(0.99, 1.05)	2.32	(1.99, 2.71)
5	0.18	(–0.12, 0.47)	1.01	(0.98, 1.05)	2.82	(2.43, 3.31)
6	0.19	(–0.18, 0.55)	1.02	(0.98, 1.06)	4.16	(3.57, 4.86)
7	0.04	(–0.40, 0.46)	1.03	(0.98, 1.08)	5.48	(4.70, 6.41)

Table 3: The RMSE, MAE, and RB of the forecasts made by the four methods, for the Southampton data. The nominal coverages of the 95% forecast intervals are also provided for the proposed model using observed and forecasted temperatures.

	1 day	2 day	3 day	4 day	5 day	6 day	7 day	Overall
	Root Mean Squared Error (RMSE)							
Persistence	6.453	6.450	6.452	6.456	6.463	6.462	6.461	6.457
Six Week Moving Average	5.094	5.094	5.095	5.091	5.096	5.098	5.108	5.097
Observed Temperatures	4.827	4.830	4.838	4.821	4.840	4.854	4.859	4.838
Forecasted Temperatures	4.832	4.841	4.923	4.994	5.034	5.221	5.268	5.019
	Mean Absolute Error (MAE)							
Persistence	5.023	5.026	5.029	5.033	5.038	5.035	5.038	5.033
Six Week Moving Average	3.961	3.961	3.960	3.958	3.962	3.965	3.972	3.963
Observed Temperatures	3.785	3.790	3.800	3.782	3.795	3.805	3.805	3.795
Forecasted Temperatures	3.780	3.797	3.845	3.915	3.949	4.056	4.117	3.923
	Relative Bias (RB)							
Persistence	0.075	0.075	0.076	0.077	0.077	0.076	0.077	0.076
Six Week Moving Average	0.079	0.078	0.079	0.080	0.079	0.0793	0.080	0.079
Observed Temperatures	0.042	0.042	0.043	0.043	0.044	0.044	0.045	0.043
Forecasted Temperatures	0.044	0.047	0.037	0.041	0.038	0.049	0.043	0.042
	Nominal Coverage (95%)							
Observed Temperatures	96.24	96.05	96.10	96.19	96.15	96.24	96.15	96.16
Forecasted Temperatures	96.05	96.43	96.06	95.58	95.36	94.94	94.48	95.57

Table 4: The RMSE, MAE, and RB of the forecasts made by the four methods, for the Cardiff data. The nominal coverages of the 95% forecast intervals are also provided for the proposed model using observed temperatures.

	1 day	2 day	3 day	4 day	5 day	6 day	7 day	Overall
	Root Mean Squared Error (RMSE)							
Persistence	7.299	7.529	7.247	7.515	7.595	7.583	7.369	7.449
Six Week Moving Average	5.722	5.934	5.858	5.989	5.827	5.986	5.715	5.863
Observed Temperatures	5.782	5.844	5.865	6.134	5.772	5.672	5.580	5.809
	Mean Absolute Error (MAE)							
Persistence	5.505	5.600	5.460	5.610	5.660	5.691	5.526	5.579
Six Week Moving Average	4.205	4.350	4.373	4.429	4.410	4.486	4.231	4.355
Observed Temperatures	4.268	4.350	4.379	4.488	4.318	4.255	4.169	4.318
	Relative Bias (RB)							
Persistence	0.051	0.077	0.074	0.075	0.105	0.085	0.059	0.075
Six Week Moving Average	0.065	0.082	0.077	0.067	0.098	0.088	0.058	0.077
Observed Temperatures	–0.028	–0.015	–0.013	–0.023	0.002	–0.008	–0.030	– 0.017
	Nominal Coverage (95%)							
Observed Temperatures	95.26	94.66	95.38	95.72	95.44	95.74	95.92	95.31

Figure 1: Boxplots showing the number of admissions by day of the week, left panel is for paediatrics, the middle panel is for adults and the right panel is for the elderly age group.

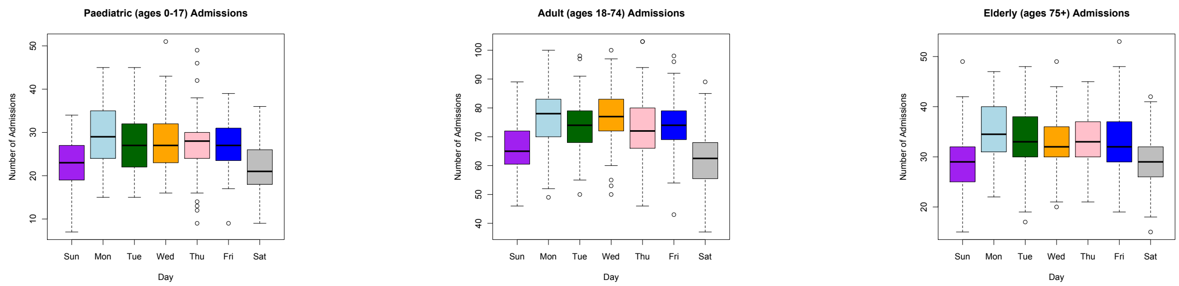

Figure 2: Multiple time series plots of number of admissions (axis on the left hand side) and temperature (axis on the right hand side) during the two years 2008-9. The left panel is for paediatrics, the middle panel is for adults and the right panel is for the elderly age group.

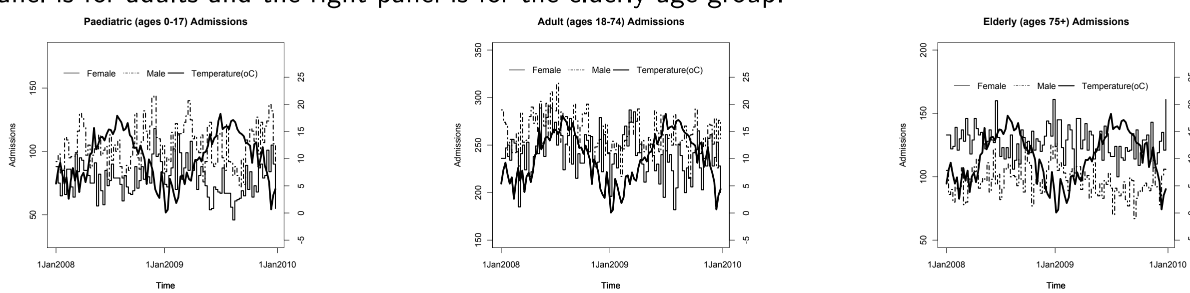

Figure 3: Boxplots contrasting the number of admissions over school holiday and non-holiday periods during 2008-9, left panel is for paediatrics, the middle panel is for adults and the right panel is for the elderly age group.

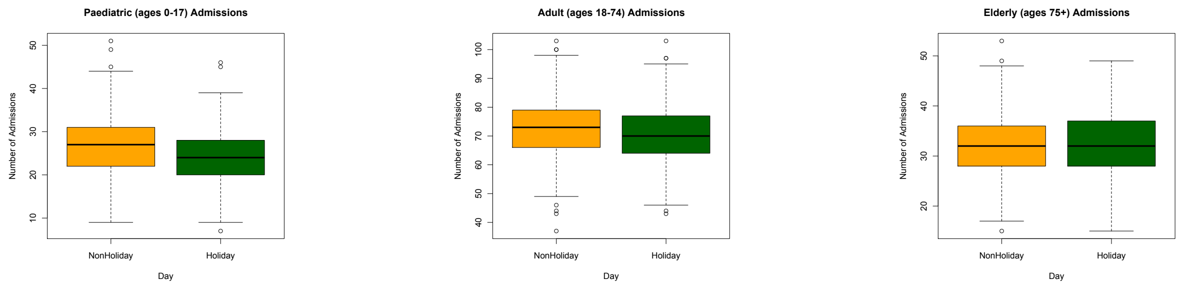

Figure 4: Boxplots showing the number of admissions by month, left panel is for paediatrics, the middle panel is for adults and the right panel is for the elderly age group. The monthly average temperatures are also plotted as solid lines in each of the plots. The axis for the temperature is on the right hand side.

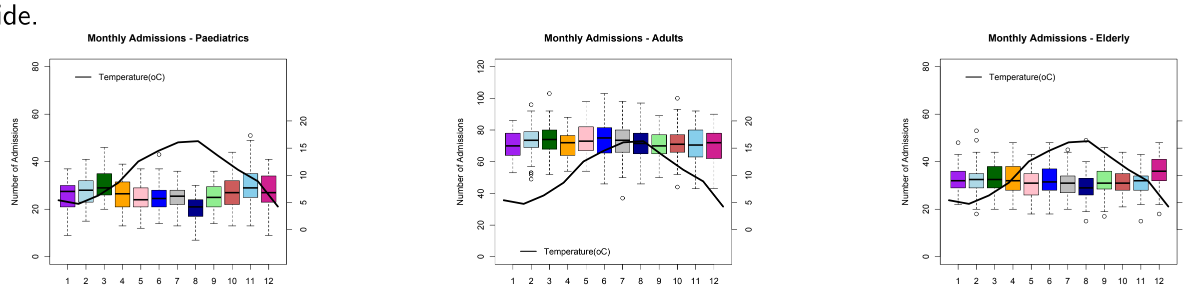

Figure 5: Plots of the residuals r_{ijt} against the lagged residuals r_{ijt-1}, for $i = \{Female, Male\}$ and $j = \{Paediatric, Adult, Elderly\}$ admissions. The top panel is for females and the bottom panel is for males.

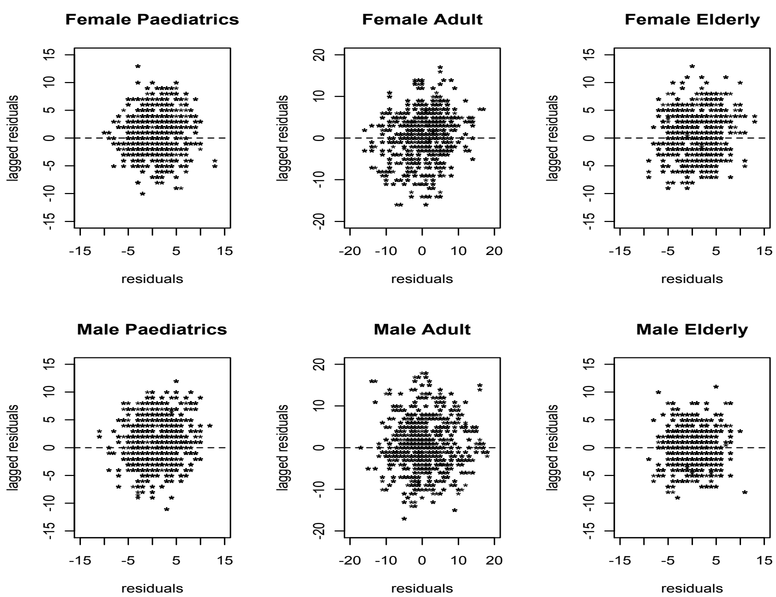

Figure 6: RMSE of three forecasting methods over the entire validation periods for the two hospitals, with mean temperature superimposed.

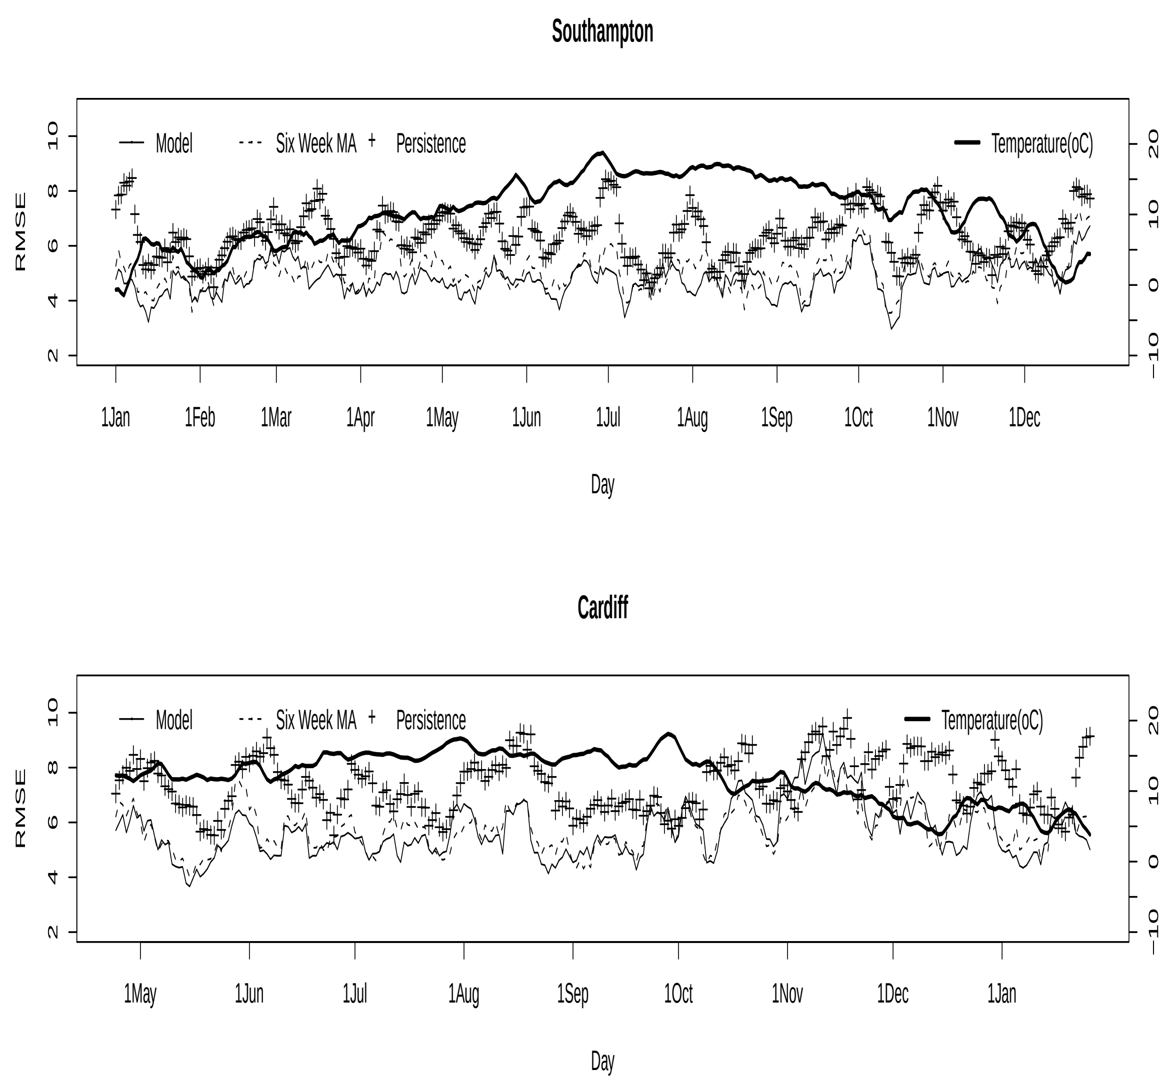

Figure 7: Boxplots of the lengths of the 50% and 90% forecast intervals.

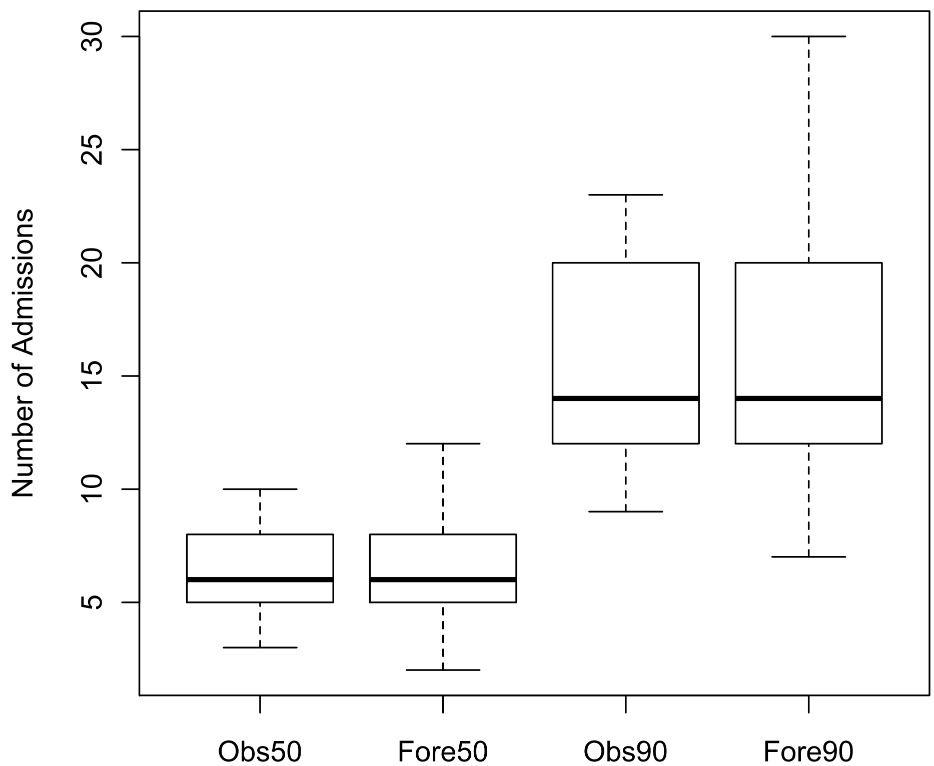

Figure 8: Lineplots showing the RMSE and the nominal coverage of the 95% forecast intervals categorised by the day of the week; the left panel is for days that fall during the school holiday periods and the right panel is for the other days. The symbol 'o' represents the forecasts using observed temperatures, and '+' represents the forecasts using the forecasted temperatures.

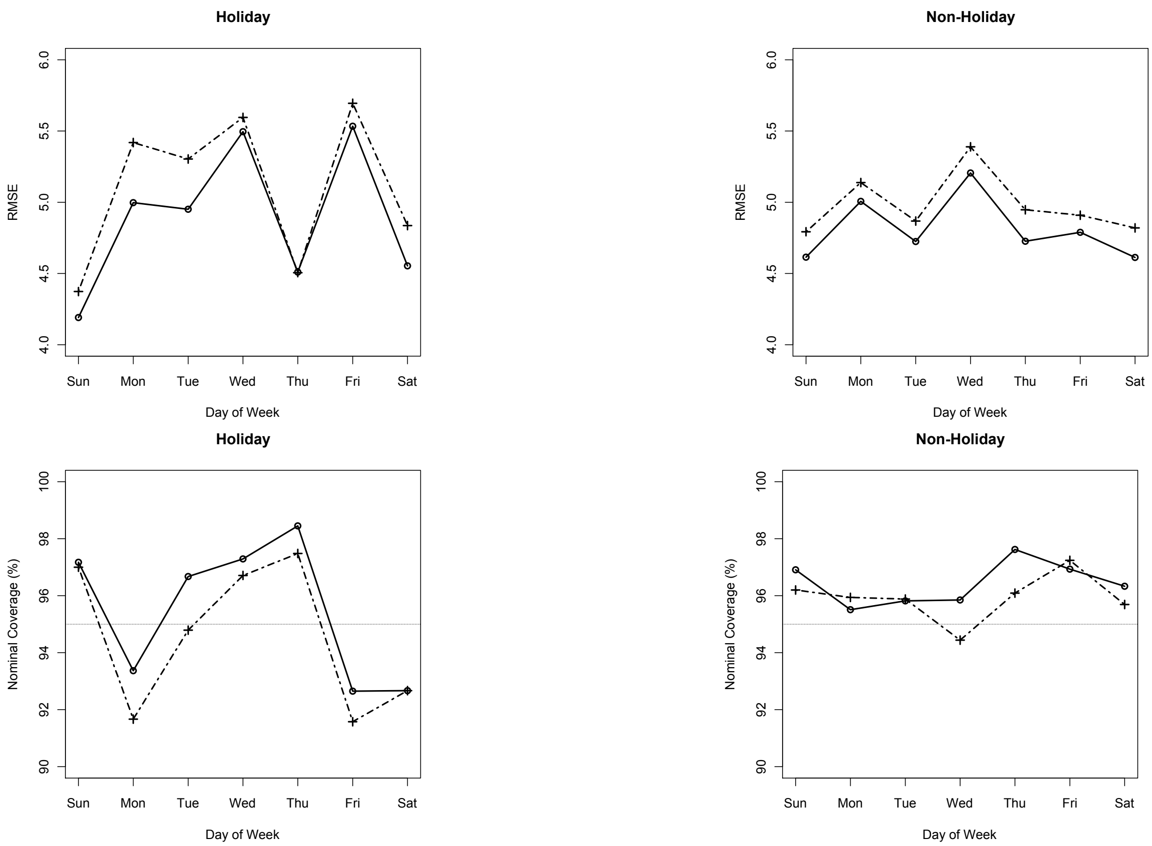

Figure 9: Lineplots showing the MAEs and the RBs categorised by the day of the week; the left panel is for days that fall during the school holiday periods and the right panel is for the other days. The symbol 'o' represents the forecasts using observed temperatures, and '+' represents the forecasts using the forecasted temperatures.