

Anticipating special events in Emergency Department forecasting

Bahman Rostami-Tabar^{a,*}, Florian Ziel^b

^a*Cardiff Business School, Cardiff*

University, Cardiff, Wales, CF24 4YX, UK

^b*University of Duisburg-Essen, Faculty of Business Administration and Economics, Germany*

Abstract

(1) Problem definition:

The Emergency Department (ED) allows patients to access a nonstop urgent medical services in the health service. A poor service delivery can often have a negative impact on the entire health service. It may put patient's life in risk and increase patient waiting time. It also increases pressure on ED seeking external resources which will consequently increase cost and cause staff sickness due to workload pressure. Moreover, it increases pressure on other health services such as Ambulance services and social care. A poor service delivery is often the outcome of a reactive strategy to a chaining demand in the ED planning system. An accurate forecast of ED attendance can be used as a planning aid towards a proactive strategy to improve the quality of service. This has now become a priority in the EDs of the National Health Service (NHS), UK. Among factors that may impact the ED attendance, special events such as public holidays and festive days play an important role. In this paper, we propose a forecasting model that outperforms benchmarks by incorporating the impact of special events and their cumulative effect.

(2) Academic / Practical Relevance:

It is crucial for emergency department to access accurate daily forecasts of attendance. Accurate forecast of daily ED attendance helps roster

*Correspondance: B. Rostami-Tabar, Cardiff Business School, Cardiff University, Cardiff, Wales, CF24 4YX, UK. Tel.: +44-

Email address: `rostami-tabarb@cardiff.ac.uk` (Bahman Rostami-Tabar)

planners in the ED to better allocate available resources. It may reduce i) cost associated with calling external resources, ii) staff sickness, iii) pressure on other health services iv) waiting time and finally improve the quality of service for major stakeholders. Since special events affect human behaviour, they may increase or decrease the demand for ED service. Most of studies that analyze ED demand: i) are descriptive in nature, ii) mainly focus on time series approaches, iii) fail reproducibility and lack rigor, (iv) do not consider modeling the type of special events, and (v) completely ignore the future uncertainty of attendance. To the best of our knowledge, there is no research that models the type of special events such as holidays and festive days to estimate daily attendance in the ED with probabilistic forecasts.

(3) Methodology:

We propose a forecasting model to generate daily forecasts of ED attendance and to analyze the impact of special events on ED demand. The model includes: i) Auto-regressive effect; ii) weekday effect; iii) long term trend effect iv) special events effects . Moreover, it provides a probabilistic forecasts to highlight the uncertainty around the forecast of ED attendance. Using a daily data set of over 6 years from a hospital in the UK, we demonstrate the effectiveness of the proposed model on forecast accuracy. We benchmark against existing approaches: 1) Naive; 2) Auto-regressive order p ; 3) Simple regression without special events effects; 4) exponential smoothing state space (ETS) model and 5) TBATS for forecasting complex seasonal time series. We analyze in-sample results to determine the effect of parameters. Moreover, we evaluate the forecast accuracy in out-of-sample using Median Absolute Error, Root Mean Squared Error for point forecasts and using Pinball score and energy score for probabilistic forecasts.

(4) Results

We show that the proposed model outperforms benchmarks across all horizons for both point and probabilistic forecasts. We provide evidences that modeling special events will improve the forecast accuracy of ED

attendance. Moreover, we describe in details how might different type of events increase or decrease the ED attendance. Results also show that our model is more robust with increasing forecasting horizon.

(5) Managerial Implications

We provide suggestions for the treatment of special events in the ED daily attendance forecasting to maximize gains in resource allocation. This is especially useful, as the incorporation of special events can improve the forecasting accuracy during public holidays and festive periods. Our model can easily be adapted to be used not only by EDs but also by other health services. It could apply to any service such as Ambulance service or hospitals trying to generate accurate forecast by considering special events. Our model could also be generalized to include more special events if necessary.

Keywords: Emergency Department, Forecast accuracy, Holidays

1. Introduction

Health Service is increasingly regarding patients as consumers of its service and aims to increase customer satisfaction; particularly in the Emergency Department (ED), which has been viewed as the shop window of the hospital service(Booth et al., 1992). This is because ED waiting time (Ang et al., 2015) can be affected by changing activities and pressures in other services such as the ambulance service, primary care, community-based care and social services. ED allows patients to access a nonstop urgent medical services. A poor service delivery can often make the difference between life and death. Figures show the number of people being treated in ED units are constantly changing (Kamali et al., 2018) and is often in excess of what the department can suitably manage(Blunt, 2014; Thomas Schneider et al., 2018). This measure of relative occupancy is an indication of the pressure in the department and naturally is associated with longer average waiting times. Increased attendees to ED and failure of patient flow through the department are caused by a multitude of factors. Some factors are internal to ED and the health service such as lack of resources, whereas others are external factors such as population growth, changes

in life style and changing behaviours due to special events such as public holidays. Special events are external deterministic events to ED that are known in advance. They can be used for modeling purposes to improve the forecast accuracy of ED attendance. Forecasting ED attendance and, therefore its workload represents a vital component of planning. The ability to forecast the number of attendance on any given day allows for more accurate staffing schedules and may reduce congestion in the ED (Xu and Chan, 2016), which consequently improve the quality of service for patients, reduces pressure on ED staff and decreases ambulance queues and reduces costs for the health system(Xu and Chan, 2016). Most of published researches in forecasting ED attendance use time series techniques, which are purely based on past historical observations of ED attendance (Marcilio et al., 2013). Few studies investigate the impact of public holidays on ED attendance (Tai et al., 2007; Sun et al., 2009; Faryar, 2013; Rotstein et al., 1997; Marcilio et al., 2013; Diehl et al., 1981; Alhusain et al., 2017). In these studies, the effect of type of holidays has not been modeled and never been used for forecasting purposes. Moreover, researches in this area fail reproducibility principles and lacks strong methodological rigor (Boylan et al., 2015; Boylan, 2016). No analysis has been conducted to accurately forecast the mid-term daily attendance in ED by looking at a comprehensive type of special events in addition to the existing effects in the historical data. This paper aims to bridge this gap by first analyzing the impact of special events on ED attendance and then proposing a model to forecast ED attendance in the presence of type of events. In addition to point forecasts, we provide probabilistic forecasts to highlight uncertainties around ED point forecasts which is crucial for decision making in the ED. The study uses a daily data set of patient attendance for 6 years from an emergency department in a major public sector acute care regional general hospital in Wales, UK. Our model can easily be adapted and it could, therefore, be useful for other EDs. It takes historical time series of ED attendance and time series of type of events as input and generates daily point and probabilistic forecasts for the required horizons. In this paper, our objectives are fourfold: (i) We examine how different type of events affect the ED attendance and provide evidence of their usefulness in

forecasting daily ED attendance; (ii) we propose a model to accurately forecast the daily ED attendance which consider different type of events, weekday effect, Auto-regressive effect, long-term trend and date effects; (iii) we provide probabilistic forecasts that quantify uncertainties in future ED attendance; and (iv) we benchmark the accuracy of our model against four time series techniques i.e. Naive, AutoRegressive, AR(p) (Box et al., 2015), exponential smoothing state space model (ETS) (Durbin and Koopman, 2012), TBATS (De Livera et al., 2011) and a regression model without considering special events as alternatives. The rest of the paper is organized as follows: Section 2 describes different category of events considered in this study. Section 3 provides an overview of the use and development of forecasting in ED and section 4 describes the proposed model and alternatives benchmarks. Section 5 describes the data set used and the setup of our evaluation, while section 6 presents the results, followed by concluding remarks in section 7.

2. Special events in forecasting ED attendance

Special events such as public holidays are specific days of the year when people might be away from their professional obligations. This may include public holidays, celebratory dates and school/university holidays, festivals and sport events. The existence of these type of events tends to change people's behaviors. ED attendance may be influenced by this change as well which consequently has an impact on the ED workload. Events shift the place of activities from the business areas to the entertainment center; i.e. while some people may stay home, others may prefer outdoor activities, travel, party and drink. These activities may increase or decrease the number of people require ED services. Thus, depending on the type of events, ED demand may increase or decrease. The occurrence of events is usually known in advance. It is the deterministic nature of an event that make it an interesting variable to be considered for forecasting. These can be modeled appropriately from a forecasting perspective. However, events occur usually once or at least once a year, so it can be seen as a rare event and rather difficult to establish a model which would be able to credit for such a day. Moreover, ED attendance may have weekly seasonal

pattern. Events may violate weekly pattern and shift the attendance to another day. In this study we consider three different type of events:

2.1. Event classification

In this study we consider three different type of events based on their characteristics: (i) fixed-date event; (ii) flexible-date event; (iii) long-date event.

Table 1: Event classification for Wales

Fixed-date	Flexible-date	Long-date
New Year Day	Pancake Day	Winter School Holiday
Valentines Day	Mothering Sunday	Spring Half-Term School Holiday
St. Davids Day	Good Friday	Spring School Holiday
St. Patricks Day	Holy Saturday	Summer Half-Term School Holiday
Halloween Day	Easter Sunday	Summer School Holiday
Guy Fawkes Night	Easter Monday	Autumn Half-Term School Holiday
Christmas Day	Early May Bank Holiday	
Boxing Day	Spring Bank Holiday	
Royal Wedding Bank Holiday	Fathers Day	
	August Bank Holiday	
	Remembrance Sunday	
	Black Friday	

The first group contains the fixed-date events. They are always occur on a fixed date each year. Typical examples are New Year Day(1 January), Christmas Day(25 December) and Boxing Day (26 December). Many national holidays linked to historically important dates of particular countries are also fall into this category. The second group contains the flexible-date events. These holidays always occur on the same weekday, however the occurrence date varies from year to year. Typical examples are Good Friday and Ester Monday. Flexible-date events may also differ depending on the geographical location. Summer bank holiday, for example, usually differ between England-Wales and Scotland. The third category of events are long-date events that contains both school and university holidays. Each school holiday term usually lasts for one week but no more than two weeks, with the exception of the Summer School Holiday which lasts for more than a month. Universities generally have three holiday terms:

(i) Winter Recess; (ii) Easter Recess; and (iii) Summer Recess. University holidays may last up to one-two and half months. School and University holidays are important to be considered for modeling the ED attendance especially in student cities. Table 2.1 provides a summary of holidays considered in this study.

3. Research background: ED forecasting

Multiple studies analyze the historical time series demand in the emergency departments. Some studies are descriptive in nature, others use time series analysis or regression to forecast the ED attendance. Rotstein et al. (1997), Sun et al. (2009) and Kam et al. (2010) analyze whether calendar seasons have an impact on the ED visits. Sun et al. (2009) show that Monday has the highest attendance in the week, however Kam et al. (2010) concludes that it is Sunday. They also observe different behaviours in winter and summer. Some studies demonstrate that public holidays increases ED patient volume (Tai et al., 2007; Zheng et al., 2007; Sun et al., 2009; DaGar et al., 2014). However, others show a decrease in ED attendance during public holidays (Faryar, 2013; Rotstein et al., 1997; Marcilio et al., 2013; Diehl et al., 1981; Alhusain et al., 2017). Baer et al. (2011) also show that local weather conditions affect the ED attendance. Diehl et al. (1981) examine the effect of calendar and meteorological factors on the ED attendance. They use a stepwise regression model to forecast the future attendance. They show that the patient flow peaked on Monday and declined steadily during the remainder of the week. they also show that less visits occurred during autumn and winter than during summer months. Moreover, they indicate that higher temperatures are associated with more visits and daytime rainfall with less. Forecasting ED attendance have widely been studied in the literature. Wargon M. (2009) provides a literature review of forecasting patient visits to ED. They state that the most used forecasting approaches are either linear regression models including calendar variables or time series models. Champion et al. (2007) use Seasonal Exponential Smoothing (SES) and Autoregressive Integrated Moving Average (ARIMA) of order (0,1,1) to forecast the monthly number of patients at the ED of a hospital in

regional Victoria. They use monthly time series from 2000 to 2005 to evaluate the forecast accuracy through Root Mean Square Error (RMSE). Results show that SES outperforms ARIMA(0,1,1). Reis and Mandl (2003) forecast both the overall visits and the respiratory-related visits using ARIMA. They show that the mean absolute percentage error (MAPE) of the ARIMA model is 9.37% for overall visits and 27.54% for respiratory visits. Using an hourly data set of ED bed occupancy from July 2005 to June 2006 for three hospitals, Schweigler et al. (2009) produce short-term forecasts of ED bed occupancy. They use three methods for each site: 1) hourly historical average, 2) seasonal autoregressive integrated moving average (ARIMA), and 3) sinusoidal with an autoregression (AR)-structured error term. They evaluate the forecast accuracy of four and twelve hours forecast in advance using root mean squared error (RMSE). They show that both the sinusoidal model with AR-structured error term and a seasonal ARIMA model predict robustly ED bed occupancy. Kam et al. (2010) use moving average (MA), univariate and multivariate seasonal auto-regressive integrated moving average (SARIMA) models to forecast daily visits of a regional ED. They use weather and calendar information as explanatory variables in building the model. The result suggests that SARIMA model provides more accurate results than MA. They suggest incorporating weather information (temperature and rain) to predict daily volumes. Marcilio et al. (2013) compare the forecast accuracy of the Generalized Linear Model (GLM), Generalized Estimating Equations (GEE), and Seasonal Autoregressive Integrated Moving Average (SARIMA) methods using total daily patient visits to an ED in Sao Paulo, Brazil, from January 1, 2008, to December 31, 2010. They use the last three months of data to measure the forecast accuracy of each model using MAPE. They conclude that GLM and GEE models provide more accurate forecasts than SARIMA model. They also state that forecast performance is better in the short-term horizon (7 days in advance) than the mid-term (30 days in advance). Calegari et al. (2016) examine the accuracy of the simple seasonal exponential smoothing (SES), seasonal multiplicative Holt-Winters (SMHW), SARIMA, and Multivariate Autoregressive Integrated Moving Average (MSARIMA) to forecast patient volumes in the ED. They consider horizons

of one, seven, fourteen, twenty-one, and thirty days. They use a daily data set of twenty-seven months for the analysis. They calculate MAPE for each forecasting horizon using ninety days as test set. They conclude that SARIMA is the most accurate approach for all horizons. Holleman et al. (1996) examine the effect of calendar and weather variables on daily unscheduled patient volume in a walk-in clinic and emergency department. They consider (i) calendar variables such as season, week of month, day of week, holidays, and federal check delivery days; and (ii) weather variables such as high temperature and snowfall as external variables. A daily data set from November 1991 to August 1994 is used for the forecasting experiment. They construct a linear regression model to predict daily patient volume. They claim that five calendar variables (season, day of week, week of month, relation to federal holiday, and federal check-delivery day) predict accurately daily patient volume based on in-sample evaluation measured by $R^2 = .86$, and $p\text{-value} = .0001$. However, they have not reported the forecast accuracy using genuine test set and relying on R^2 could be misleading as a good fit does not guarantee accurate forecasts. Batal et al. (2001) examine the impact of the calendar variables in staff scheduling using a daily patient volume from February 1998 to July 2000. They show that when the calendar variable is the only variable in the model, the model shows an accuracy of 73%. Accuracy measures for the validation set are not reported in the study. Sun et al. (2009) use SARIMA model to forecast daily patient volume for each patient acuity level. The data set includes daily patient attendances at ED between July 2005 and March 2008. The last six month is used as test and the rest for training the model. MAPE is used to choose the best-fit model. They fit separate ARIMA models to the three categories of acuity and overall data. The authors conclude that the predictions had a good accuracy. They observe that the impact of weather is not significant. They claim that ARIMA model is effective for both short-term forecasts (weekly) and long term (three months). The author do not consider any benchmark method in the study. The literature review reveals some limitations in forecasting ED attendance which will be summarized as follows: (i) the effect of type of events is not modeled and never been used for forecasting purposes. They are mainly descriptive and if used for

forecasting, it does not differentiate the type of events; (ii) some studies in this area lack rigour of experimental design, i.e. there is no benchmark method or forecast accuracy is not reported; (iii) most of studies are not reproducible, it is almost impossible to reproduce the forecasting models and results; (iv) they are limited in terms of the length of historical data used for the training purpose and forecast performance evaluation and finally (v) none of them consider uncertainty of future forecasts which provides the whole picture of future ED attendance for risk management. No analysis has been conducted to model the type of events in addition to the calendar effects and existing patterns in historical data to accurately forecast the ED daily attendance. Additionally, to the best of our knowledge, no study looks at the probabilistic forecasts in the ED forecasting that provides more information associated with point forecasts that inform decision makers about uncertainties of future ED attendance. Considering these limitation from the literature, we propose a forecasting model to fill these gaps.

4. Proposed model

We consider a high-dimensional time series model where the mentioned components might be zero after applying a certain shrinkage estimator to fit the model. It has several components which allow possible interpretations, including a) Weekly effects; b) Date based effects (= annual effects); c) Event effects; d) Long-term trend effects; e) Autoregressive effects; and f) Weekly autoregressive effects. The weekly effects a) cover deterministic weekly variations, especially potential weekend effects. Date based effects b) cover annual effects that occur always at a certain date or a period of dates. Thus, these effects repeat every year, though there are not strictly periodic as a calendar year has either 365 or 366 years. These effects include automatically fixed-date event effects. Other event effects such as flexible-date and long-date are covered by c). The long-term trend component d) shall take care about long term structural changes, e.g. induced from a different life style or a change in the population of the catchment area. The autoregressive components e) and f) cover the autoregressive effects. Whereas the first one covers standard linear relationships

to the past of the time series, the latter one covers periodically autoregressive effects. They can be understood as a periodically time-varying autoregressive effects with a period of 7 which corresponds to a week. Thus, the autoregressive effect on a Sunday might be different to this one on a Monday.

We denote Y_t as the number of patients in attendance at day t . To define the model, we require some dummies, which are presented as following:

- DoW_t^k for the k -th day-of-the-week dummy. I.e. DoW_t^1 is 1 if t is on a Monday and 0 otherwise, up to DoW_t^7 is 1 if t is on a Sunday and 0 otherwise.
- DoY_t^k for the k -th day-of-the-year dummy. I.e. DoY_t^k is 1 if t the k -th day of the year if it has 365 days and 0 otherwise. If a t is in a leap year (which has 366 days) the 60th day (the 29th February) is ignored from the day counting. Itself it takes the same value as the 28th of February and bot days will have a joint dummy variable.
- E_t^k is 0 if t is at the k -th event or in the k -th event period. We consider 12 single events which are all flexible-date events and 6 long event periods which are school and University holidays (see Table 2.1). Additionally, we take special care on the the impact of Spring School Holidays. This holiday period is usually around Easter - either Easter is after the school holidays, it is before the school holidays, or it is in the middle of the school holidays. Hence, we separate it into three sub-periods: a dummies for the period before Good Friday, a dummy for after the Easter Monday, and a dummy for the four day period from Good Friday to Easter Monday.
- BS_t^k the k -th B-spline basis of degree 12 with 21 knots and the full support (including the in-sample data and the forecasting horizon), see e.g. Ziel and Liu (2016). The choice of the degree and the number of knots seems somehow arbitrary. In fact, the choice with a degree of e.g. 13 would lead to very similar results. However, the important aspect for the long-term trend basis functions is that their support covers more than a year to avoid the interactions with annual effects or short-term effects.

For the 4 dummy/basis functions DoW_t^k , DoY_t^k and BS_t^k we also define the cumulative version across the basis functions, so $\text{DoW}_t^{\text{cum},k} = \sum_{j \leq k} \text{DoW}_t^j$, $\text{DoY}_t^{\text{cum},k} = \sum_{j \leq k} \text{DoY}_t^j$ and $\text{BS}_t^{\text{cum},k} = \sum_{j \leq k} \text{BS}_t^j$.

The idea behind the cumulative basis function is to model the mentioned effects in a different way, due to another parametrization. For instance, DoY_t^k models the day-of-the-year effect in an absolute manner. So if t is the k -th day of the year then there is an absolute effect of this day, like a shock that is active for this particular day. This is highly relevant for special events. In contrast, $\text{DoY}_t^{\text{cum},k}$ models not the absolute effect of the particular day of the year but the change. So if this dummy is active then the effects persists for the remaining basis function period, so a year for $\text{DoY}_t^{\text{cum},k}$. This is particularly important to model special events effects that last multiple days, like summer or winter holidays which potentially last multiple weeks.

Following the description of the elements of the model, the proposed model is given by

$$\begin{aligned}
Y_t = & \beta_0 + \underbrace{\sum_{k=1}^7 \beta_k \text{DoW}_t^k + \beta_{7+k} \text{DoW}_t^{\text{cum},k}}_{\text{Weekday effects}} + \underbrace{\sum_{k=1}^{365} \beta_{14+k} \text{DoY}_t^k + \beta_{379+k} \text{DoY}_t^{\text{cum},k}}_{\text{Fixed-date effects}} \\
& + \underbrace{\sum_{k=1}^{22} \beta_{744+k} \text{E}_t^k}_{\text{Flexible-date and long-date effects}} + \underbrace{\sum_{k=1}^{22} \beta_{766+k} \text{BS}_t^k + \beta_{788+k} \text{BS}_t^{\text{cum},k}}_{\text{Long-term effects}} \\
& + \underbrace{\sum_{k=1}^{28} \beta_{810+k} Y_{t-k}}_{\text{Autoregressive effects}} + \underbrace{\sum_{k=1}^{28} \sum_{j=1}^7 \beta_{838+28(i-1)+k} \text{DoW}_t^j Y_{t-k}}_{\text{Weekly autoregressive effects}} + \varepsilon_t \tag{1}
\end{aligned}$$

4.1. Parameter estimation

The model has in total 1035 parameters which potentially describe the overall behaviour. However, not all of these parameters are active in the final model used for forecasting. A parameter is selected if it is significantly different from zero. This selection procedure is done by an estimation algorithm. We consider a shrinkage based least square approach, namely the elastic net, for the estimation (Hastie et al., 2015). To define the elastic net, we note that the model (1) can be represented as a standard linear model given by representation

$$Y_t = \boldsymbol{\beta}' \mathbf{X}_t + \varepsilon_t$$

with parameter vector $\boldsymbol{\beta}' = (\beta_0, \beta_1, \dots, \beta_{1034})'$. Then, we consider the observable period $\mathbf{Y} = Y_1, \dots, Y_T$ with the corresponding regression matrix $\boldsymbol{\mathcal{X}} = (\mathbf{X}'_1, \dots, \mathbf{X}'_T)'$. Moreover, we scale both \mathbf{Y} and $\boldsymbol{\mathcal{X}}$ such that each row has mean zero and standard deviation one, denoted by $\tilde{\mathbf{Y}}$ and $\tilde{\boldsymbol{\mathcal{X}}}$. The symbol $'$ is used for transpose.

$$\hat{\tilde{\boldsymbol{\beta}}} = \arg \min_{\boldsymbol{\beta}} \|\tilde{\mathbf{Y}} - \boldsymbol{\beta}' \tilde{\boldsymbol{\mathcal{X}}}\|_2 + \lambda \left(\alpha \|\boldsymbol{\beta}\|_1 + \frac{1 - \alpha}{2} \|\boldsymbol{\beta}\|_2 \right) \quad (2)$$

is the scaled elastic net estimator which depends on two tuning parameters $\lambda \geq 0$ and $\alpha \geq 0$. Note that we can obtain the unscaled lasso estimator $\hat{\boldsymbol{\beta}}$ simply by rescaling $\hat{\tilde{\boldsymbol{\beta}}}$.

We choose $\alpha = 0.5$ as it yields accurate forecasting results in a forecasting study Uniejewski et al. (2016). Moreover, it reduces the computational cost of finding an optimal α . As $\alpha = 1$ gives the popular lasso estimator, and $\alpha = 0$ the ridge estimator, the choice $\alpha = .5$ seems to be a intuitive compromise. Note that the elastic net for $\alpha < 1$ retains the sparsity property as the lasso does. Therefore, the more irrelevant a parameter is the more is it shrunk towards zero, including the special case of exactly zero. The shrinking effect depends on the tuning parameter λ . The larger λ the more parameters will be zero. For $\lambda = 0$, we receive the OLS solution as a special case where all parameters are included in the model.

We optimize the tuning parameter λ by minimizing the Hannan-Quinn-Information criterion (HQC) on an exponential grid. The HQC lead to overall robust results in a large forecasting study in Ziel and Weron (2018) across several data sets in the energy forecasting context. The HQC is a special choice of the generalised information criterion (GIC)

$$\text{GIC}(\kappa; \boldsymbol{\beta}) = T \log(\text{RSS}(\boldsymbol{\beta})/T) + \kappa K(\boldsymbol{\beta})$$

with RSS as residuals sums of squares, κ as penalty parameter and $K(\boldsymbol{\beta})$ as number of non-zero parameters in the considered model. The penalty parameter κ has to be chosen by the modeller. The HQC follows from the choice $\kappa = 2 \log(\log(T))$. It can be regarded as a compromise between the popular Akaike information criterion (AIC) followed by $\kappa = 2$ and the conservative Bayesian information criterion (BIC) with $\kappa = \log(T)$.

We also examined the tuning of the λ parameter via block-wise cross-validation. But the results did not improve compared to the information criterion based approach but it is computationally more demanding.

5. Empirical evaluation

5.1. Data

Data used in the study was counts of patients arrival times at ED between January 2010 and March 2016, extracted from the ED administrative database. We aggregate the patient arrival times to obtain the daily attendance between 01 January 2010 and 31 March 2016 that is used for empirical evaluation in this study. A sample of 3 years of data is given in Figure 1.

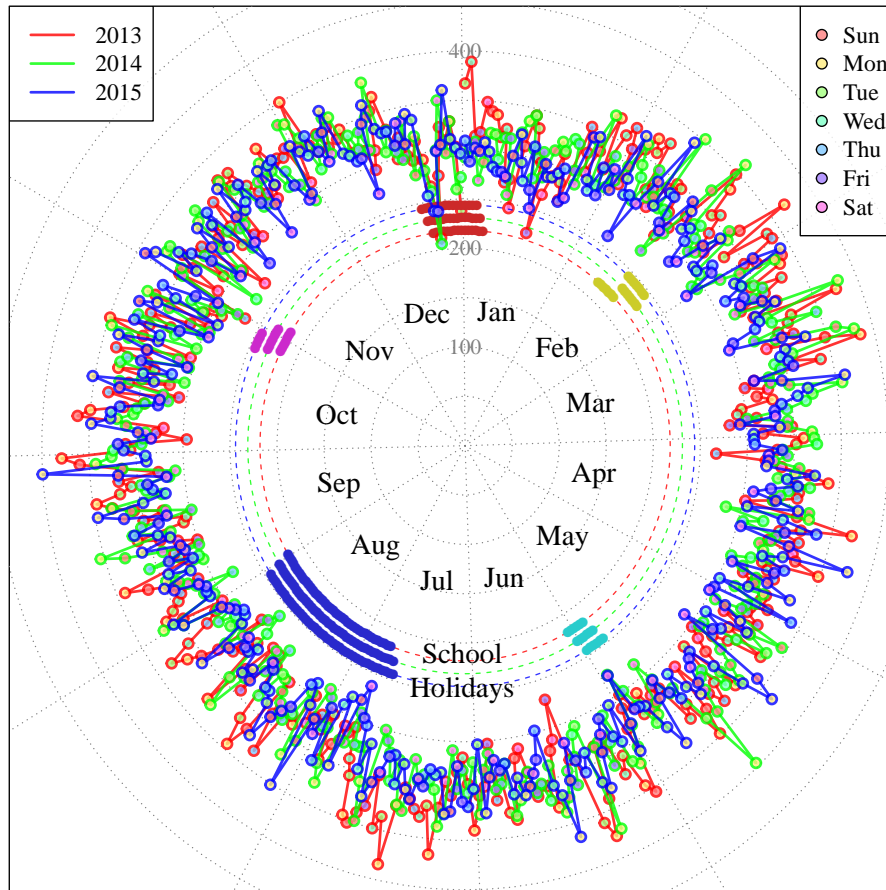


Figure 1: Daily ED attendance of three selected years (2013, 2014, 2015) in polar coordinates.

The figures illustrates the data in polar coordinates such that a full circle corresponds to a year and the distance to the origin matches the number of

patient arrivals. Although, data is noisy but some systematic structures are visible. The day-of-week effect is highlighted by different fill colors. It is clear that Mondays (yellow fill) have the peak followed by Sundays (pink fill). We have also highlighted long-date holidays such as school holidays for three years. we can observe the decreasing effect of long-date holidays. This will be confirmed later where we discussed the effete of events in the section 6.

5.2. Benchmark models

We consider five benchmark models to compare the performance of the proposed model. They include a naive benchmark model, a more sophisticated autoregressive model, an exponential smoothing model, a model for forecasting complex seasonal time series and a model with similar design to the proposed one to highlight the impact of special events. These models are described in this section.

1. Naive model:

The naive model is a very simple model, that corresponds to historic simulation. Thus, we have the assumption that there is no structure in the ED arrivals and the observed past continues to the future. Formally, this model is given by $Y_t \sim F$ where F is the unknown distribution that we estimate by the empirical distribution function (ecdf) of the given sample.

2. AR(p):

The second benchmark model is an AR(p) with p selected on the grid $1, \dots, p_{max} = 365$. Formally this is

$$Y_t = \phi_0 + \sum_{k=1}^p \phi_k Y_{t-k} + \epsilon_t \quad (3)$$

where ϕ_k denote the autoregressive parameters, ϕ_0 the intercept and ϵ_t the error term. We select the optimal p by minimizing the AIC (Akaike information criterion). For simulating from the process, we apply residual based bootstrap as for the proposed method above. Note that every ARIMA process (also SARIMA) can be rewritten as an AR(∞). As every AR process has an exponentially fast decaying memory, it can be well

approximated by a high order AR process. Thus, many of the aforementioned ARIMA-type models can be regarded as nested into the considered model.

3. Exponential smoothing:

The exponential smoothing model is incorporated using the corresponding implementation of the forecast package in R. We use the `ets()` function in forecast package to generate daily forecasts, see Hyndman and Khandakar (2008). Note that the structure identification is done automatically, but the input data is assigned with a weakly seasonality of 7. Moreover, in all cases of the forecasting study the structure is identified so that a simple exponential smoothing with additive errors is optimal. Note that in contrast Calegari et al. (2016) applied a multiplicative Holt-Winters approach.

4. TBATS:

De Livera et al. (2011) proposed a model to deal with time series exhibiting multiple complex seasonalities. TBATS includes a Box-Cox Transformation, ARMA model for residuals and a trigonometric expression of seasonality terms. The later one not only gives the model more flexibility to deal with complex seasonality but also reduces the parameters of model when the frequencies of seasonalities are high. We fit a TBATS model to our daily time series using the `tbats()` function in the forecast package of R Hyndman and Khandakar (2008).

5. Proposed model without date and holiday effects:

Here, we simply apply the same model and estimation method as for the proposed model (1), but remove the term for the date and holiday effects.

5.3. Probabilistic forecasting setup

In this study, we provide probabilistic forecasts in addition to point forecasts. It quantifies uncertainties around generated forecasts that will allow planners in ED to better manage risks in operational planning. Figure 2. illustrates an example of probabilistic forecast generated by the proposed model. The 6 years and 3 months of data were used in a rolling window forecasting study with re-estimation. We consider always 4 years of in-sample data and forecast the next




Figure 2: Example, prob. forecast

$H = 42$ which corresponds to 6 weeks and is relevant for operational planning. Formally, we have Y_1, \dots, Y_T (with $T = 4 \times 365 + 1 = 1461$) as in-sample data and want to forecast Y_{T+1}, \dots, Y_{T+H} . Thus, we are interested in the full H -dimensional probability distribution of the random vector $(Y_{T+1}, \dots, Y_{T+H})$.

To tackle this sophisticated problem we apply ensemble forecasting methods. Hence, we report a large ensemble of $M = 5000$ simulated paths. These paths are simulated from the estimated model using residual based bootstrap recursively. The large ensemble is used as approximation of the forecasting distribution of interest. Every characteristic (e.g. mean, median, variance) can be computed from based on sample statistics. We also use this property for the evaluation.

5.4. Evaluation scheme and metrics

We consider four different evaluation measures, two point forecasting metrics and two probabilistic metrics. Forecasting horizon is chosen to be $H = 42$ which corresponds to 6 weeks that is relevant for operational planning in the EDs of NHS Wales. It may vary within and across EDs between 7 days (1 week) and 42 days (6 weeks).

The point forecasting metrics are the MAE (median absolute error) and RMSE (root mean square error) for point forecasts (Hyndman and Koehler, 2006). The first one is strictly proper for the median, the latter one for the mean. Hence, the MAE is minimized only for the optimal median forecast, and the RMSE only for the optimal mean forecast. Thus, they can identify the best

median and mean forecasting model. To get the median and mean from a given ensemble forecast, we evaluate the corresponding sample median and sample mean of the ensemble. As we perform a H -step ahead forecasts we evaluate the MAE and RMSE for each forecasting step separately.

Additionally to the standard point forecasting measures we apply two probabilistic forecasting measures: the pinball score (also known as quantile loss) such as the energy score.

The pinball score is a strictly proper evaluation criterion for quantiles. We evaluate it on a dense probability grid, for all percentiles (1%, ..., 99%) as popular in forecasting (Hong et al., 2016). As such dense grid approximates the underlying H marginal distributions of (Y_1, \dots, Y_H) well it can be regarded as probabilistic evaluation measure. If the distance of the considered probability grid (here 1%) is taken to be zero (as a limit) then the average pinball score across the grid converges to the CRPS (continuous ranked probability score), which is strictly proper for univariate distributions. Similarly as for the MAE and RMSE, we can evaluate the pinball score for each forecasting horizon, but also across the 99 percentiles.

Finally, we consider the energy score, see Gneiting and Raftery (2007). In contrast to the aggregated pinball score across the percentiles or the CRPS, this is a strictly proper evaluation criterion for the full H -dimensional multivariate distribution. Therefore it is suitable for detecting misspecification in the marginal distribution structure of the forecast such as dependency structure. Due to the high complexity of the energy score, it is not very popular in forecasting so far, but in the last few years more and more applications showed up, especially in weather and energy forecasting.

6. Results and discussion

6.1. Evaluate the effect of variables

Prior to the evaluation of all effects, we have a look at the residuals of the fitted model to determine whether the proposed model captured all genuine patterns presented in the time series. In Figure 3, we observe the sample autocorrelation plot (acf) of the residuals and the squared residuals, such as the

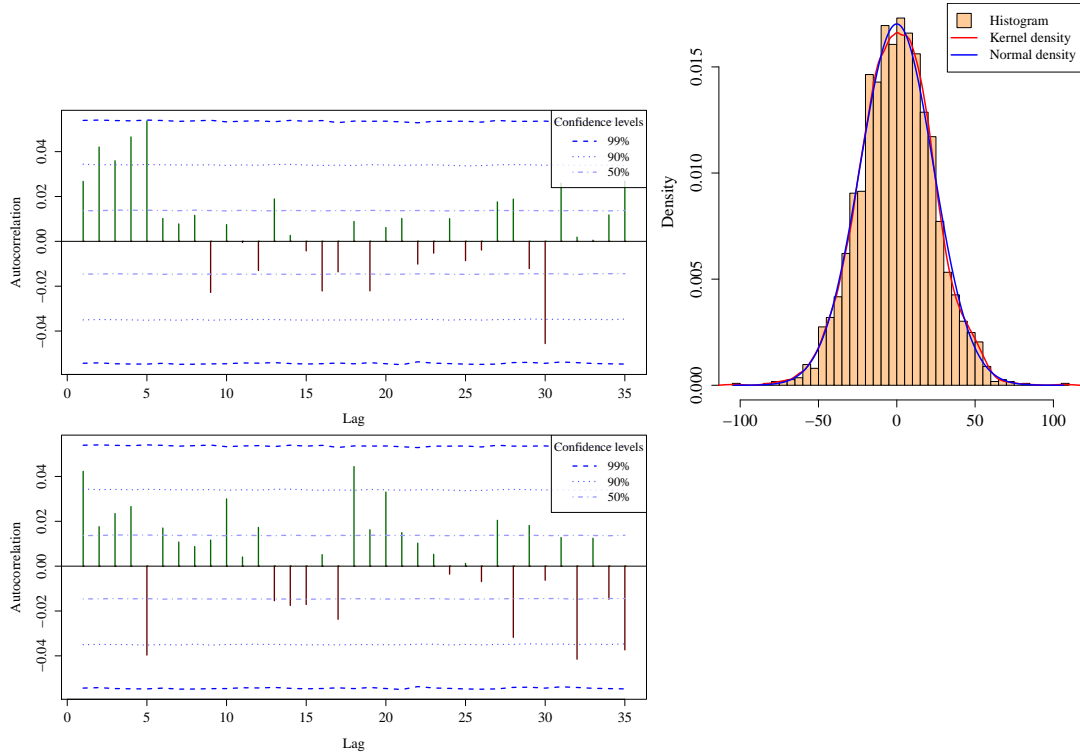


Figure 3: Sample autocorrelation function of $\hat{\varepsilon}_t$ and $\hat{\varepsilon}_t^2$, with histogram of $\hat{\varepsilon}_t$

histogram of the residuals. Both autocorrelation plots show no series correlation. Thus, the residuals seem to be uncorrelated and homoscedastic. We can see from the histogram that the residuals have a unimodal distribution without heavy tails. The density is close to a normal distribution with correspondingly matching variance. Still, note that for the estimation procedure the normality assumption is not required for an OLS estimation method. However, if the residuals follow a normal distribution, it would underline the suitability of the estimation procedure as the normally distributed likelihood approach coincides to the OLS. It is important to note that residuals are well behaved. Thus, interpretations of the model output are valid, even in-sample. Moreover, as the autocorrelation functions does not show series correlation, we can conclude that basically all autoregressive and deterministic information is covered. No substantial improvements can be expected unless external regressors are considered.

6.2. Impact of variables and interpretation

Figure 4 shows the model fit (conditional mean) with the corresponding unconditional mean estimate. We see clearly annual pattern in the data, especially a strong the winter holiday in December/January dip and a smaller summer holiday period. If you look carefully at the graph you also observe weekly variations. So it seems that weekly effects are captured as well.

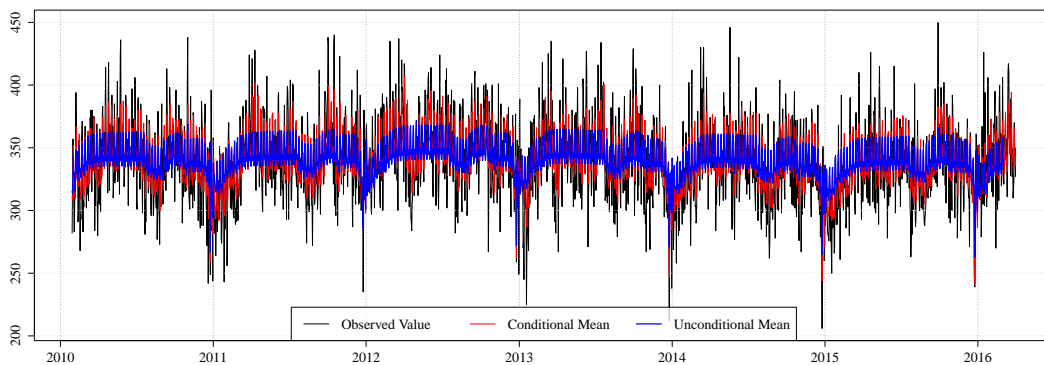


Figure 4: Observed values of Y_t with estimates of the mean $\mathbb{E}(Y_t)$ and conditional mean $\mathbb{E}(Y_t|Y_{t-1}, Y_{t-2}, \dots)$.

In Figure 5, the fitted long-term trend is visualised. We observe that there are some long-trend components estimated. However, in absolute terms this effects are very small, and barely visible in Figure 4.




Figure 5: Weekly effects and long-term trend effect.

Figure 6 presents the autoregressive effects, including the weekly periodically autoregressive effects. We see that the major effects are driven by the non-




Figure 6: Estimated autoregressive parameters, including periodically autoregressive effects

periodic parameters. Here, the most recent value (lag=1) has the strongest impact. In the memory, we also observe the weekly structure, so lag multiples of 7 (14, 21, 28) show some impacts. Next to the standard autoregressive effects, we observe some periodic effects as well. We see that the short term memory gets especially adjusted for Mondays and Sundays. These effects, are the only explanation for actual weekly seasonal structure in Figure 4. Thus, we see that even those small autoregressive parameters can have a substantial impact on the model performance. Here, the authors want to add that if the weekly periodically autoregressive parameters are removed from the model, then next to a small reduction in the model performance, the weekday dummies 5 show positive effects.

Finally, we discuss the effect of special events. We start by looking at the impact of fixed-date events.

Figure 7 depicted the effects of the date based dummies over a year with highlighted fixed-date events. We observe that there is a reduction in the baseline intensity during the colder periods of the year (end of October to end of February) where less patients visit the Emergency Department. Around the Christmas/New Years period (mid December to mid January) there seems to be an extra reduction in the effect. Moreover, we see that not all fixed-date events have an effect of the hospital attendance. Still, we have some interesting pattern.

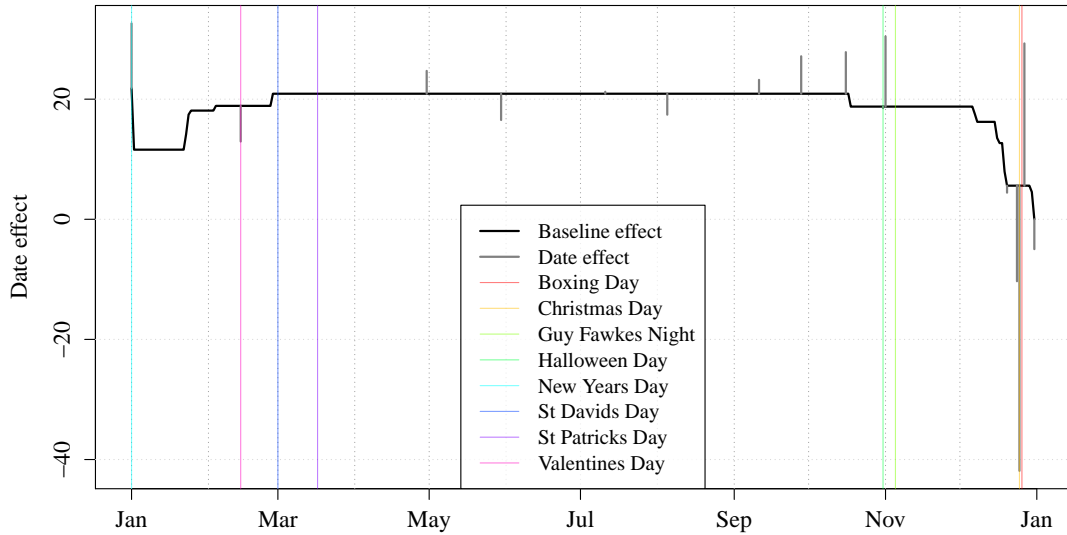


Figure 7: Estimated date based effects with highlighted fix-date events.

We observe strong impacts of Christmas, which is not surprising. We have a clear reduction on Christmas Eve (24 Dec) and Christmas Day (25 Day) but in contrast a strong positive peak on 27 Dec., the day after boxing day (26 Dec.). That might be attributed to the strong reduction around Christmas, that ill people go to hospital after the Christmas feast. Furthermore, we see a reducing effect on New Years Eve (31 Dec) but a strong increase in the attendance on New Years Day (1st Dec). A similar effect can be observed at Halloween (31 Oct). On Halloween we have a small reduction in the attendance but followed by a large increase on the next day. Likely some party people celebrated 'too much' over night and end up in the hospital the consecutive day. But there are other examples of holiday effects. For instance, at Valentines (14 Feb), there is a significant reduction the hospital visits. Maybe, directly love related events lead to less violence in everyday life and thus less hospital visits. Moreover, there are some date effects which are clearly visible and not linked to a public holiday, e.g. there is an effect on 30 April and 30 May.

Figure 8 illustrates the effect of flexible-date and long-date events. We see that most of the events in these categories have no significant impact. However, if there is an effect it is a reduction effect. We see more impact of holiday periods, namely of Autumn Half-Term Holiday, Summer Holidays. The latter one has a stronger reduction impact. Moreover, the Summer Holiday dummy is the main reason for the visible summer reduction effect visible in Figure 4.

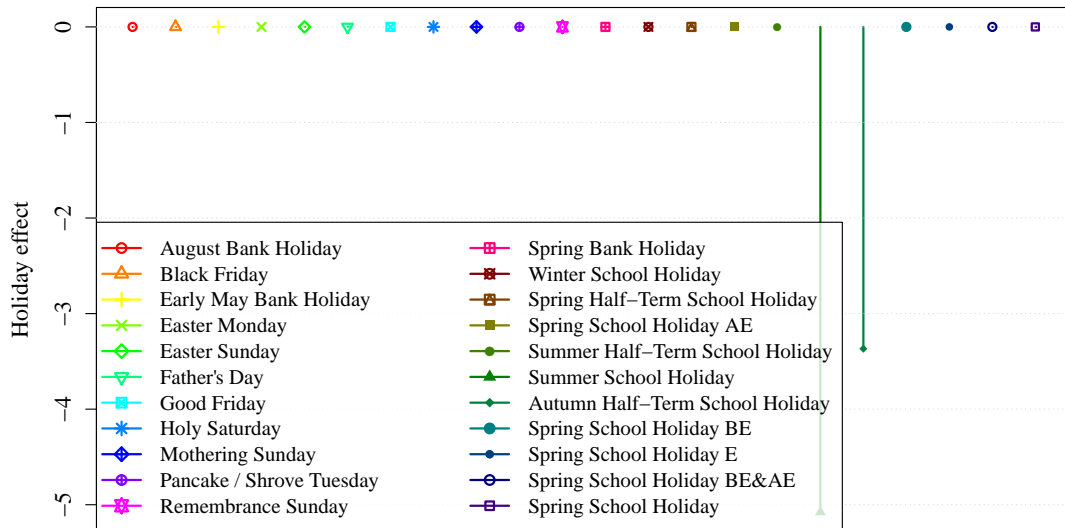


Figure 8: Estimated flexible-date and long-date effects.

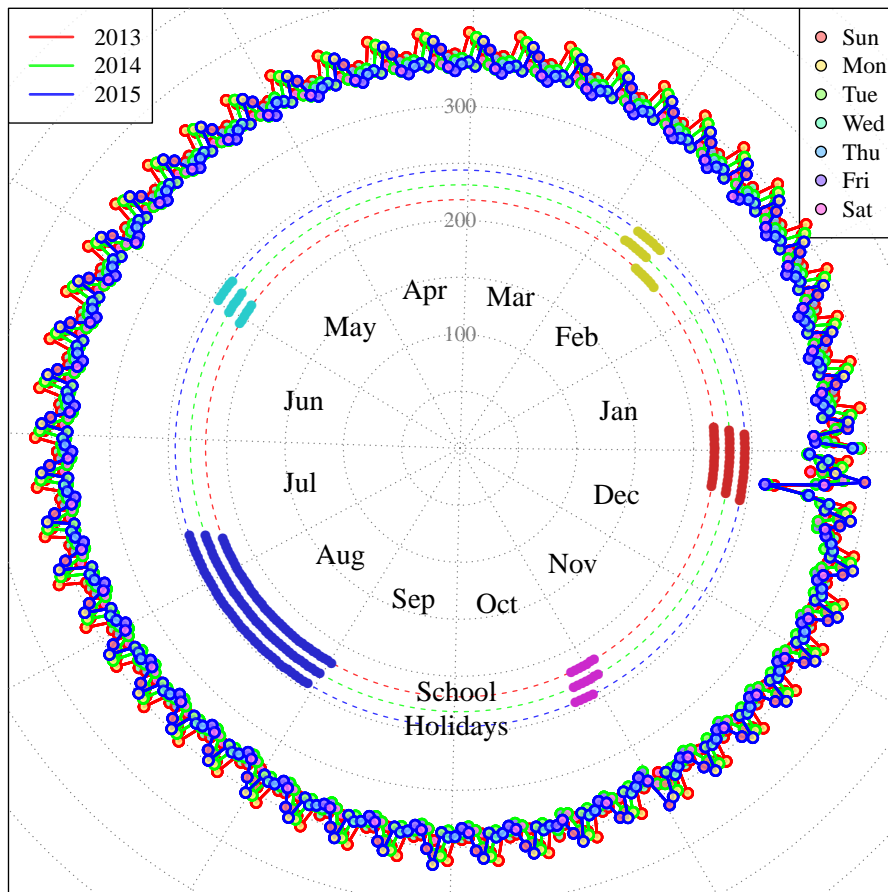


Figure 9: Estimated expected value of the daily hospital arrivals across three selected years in polar coordinates.

Obviously the the effect of School holidays and flexible-date holidays vary from year to year as the holidays shift. This effect overlaps with the general

effect that weekdays shift from year to year. It means that if 1 January is on Monday in a certain year, it will be on Tuesday or Wednesday next year, if the recent year is a leap year. The overall behavior of the model is illustrated in Figure 9 based on the structured captured in the model where noise has been filtered out. Here, the estimated expected value of the arrivals is visualized in polar coordinates with annual cycles. The distance from the center matches the corresponding number of arrivals. The mentioned weekday shift is clearly visible. Additionally, we see the strong Christmas effect, where the expected arrivals drop clearly below 300 whereas it is clearly above 300 for the remaining year. We can also observe that the effect of long-date events are captured.

6.3. Forecast accuracy evaluation

In this section, we evaluate the forecast accuracy of our model and its benchmarks using Out-of-sample data. First, we focus on the point forecasting results. Figure 10 shows the MAE and RMSE values across the forecasting horizon for all considered five models. We observe that overall, the forecasting behaviour




Figure 10: MAE and RMSE of considered models across the forecasting horizon.

concerning the MAE and RMSE are relatively similar. This is intuitive because the structure of errors looks relatively symmetric (see Fig. 4) which yields similar median and mean forecasts. Overall, we observe that the proposed model outperforms all benchmarks consistently for all forecasting horizons. Especially for longer forecasting horizons the proposed model remains a high forecasting accuracy compared to other methods. In general, the AR(p) and TBATS are relatively competitive benchmarks compared. Further, we see that the impact of special events effects in the proposed model is especially important for the longer forecasting horizons. Whereas, the proposed model without special events effects reduces its forecasting performance consistently. The proposed model not only remains accurate across the forecasting horizon but also it is clear that our model is robust against increasing horizons than the other methods.




Figure 11: Pinball score across the forecasting horizon (top) and across the quantile levels (bottom) of all considered models.

In evaluating the probabilistic forecast accuracy of our model and its benchmark, we first focus on the marginal distributional characteristics evaluated

by the corresponding pinball score. Figure 11(top) shows the corresponding pinball score aggregated across the quantile grid for each forecasting horizon. Figure 11(bottom) presents pinball score aggregated across forecasting horizons for each quantile grid.

Results show that the proposed model captures the density structure for all horizons consistently better than the benchmark models. When studying the results concerning the accuracy for single quantiles, we observe that the major part for the improvement comes from the centre quantiles. For the extremes (e.g. 1% and 99% quantile) the proposed models has no clear improvement in the forecasting accuracy compared to other models.




Figure 12: Box plot of energy score for all considered models.

Finally, we study the results concerning the fully multivariate forecasted distribution evaluated by the energy score given in Figure 12. Obviously, we have only one energy score for the full horizon. Therefore the box plot of all individuals scores across the forecasting horizon is chosen for illustration. We see that all overall the proposed model has the best energy score. It seems that main improvement of the accuracy compared to the proposed model without holiday and date effects goes back to the reduction of the upper tail. For instance there is no single day with an energy score above 200 whereas there are a couple of these days for the remaining models.

Our results show that the proposed model not only outperforms the benchmarks in terms of generating the correct value of future ED attendance but also it is the most accurate one when it comes to providing consistent information about the uncertainty around the ED attendance. It provides the entire probability distribution for the ED attendance for a given day. Our model offers

the ability to capture a range of possibilities of ED attendance for any given day that is not contained in the point forecasts. This is very important for decision makers and planners in EDs as it may help them to assess the risk and make better decisions in operational planning. Further research attempts are required on how to use probabilistic forecasts to inform decisions in the emergency department and on how to interpret output which are useful.

7. Concluding remarks

Accurate forecasts of ED attendance can help roster planners to better manage the risk of over and under capacity resource allocation. This is crucial for the health service management as both cases have a huge impact on all health service stakeholders such as patients, ED itself, ambulance services and social care services, etc. They may increase cost, increase staff sickness, waste resources, put pressure on other health services and consequently deteriorate the efficiency of the entire health service. This paper proposes a forecasting model for daily attendance in emergency departments. The model is high dimensional and includes many potential predictor variables which are classified into : i) Auto-regressive effect; ii) weekday effect; iii) long term trend effect v) special events effects.

Not all potential predictor variables are present in the forecasting model. We use an estimation procedure that decide which parameter is significantly different from zero that is kept in the model and will be used to analyze the impact of selected predictors on ED attendance and to generate the daily forecast of ED attendance. In generating both point and probabilistic forecasts, we compare the effectiveness of our model in terms of forecasting performance against four time series method and a regression model without accounting for the effect of spacial events. The benchmarks are: 1) Naive; 2) Auto-regressive order p ; 3) exponential smoothing state space (ETS) model; 4) TBATS for forecasting complex seasonal time series and 5) A regression model without special events effects. In summary, we believe that the problem setting we have considered is a very realistic one. We use daily attendance data from a major ED in the UK for a period of over 6 years. Four different accuracy measures are used to

evaluate the forecast accuracy performance of all models in generating a daily forecast, 42 days in advance. The main findings of this study can be summarized as follows:

- We propose a forecasting model that outperforms all benchmarks in generating both points and probabilistic forecasts for daily ED attendance. To the best of our knowledge, probabilistic forecasts have never been used to forecast ED attendance. Although this needs more investigation to be linked to operational planning but it provides more information about uncertainty of future ED attendance which can be used for risk management in allocating resources for staffing.
- We provide evidences that incorporating special events in the forecasting model improves the forecast accuracy. We strongly recommend decision makers and planners to take into account the increasing or decreasing impact of special events in taking decisions.
- We observe that the following fixed-date events have an increasing effect on AE demand : 1st of January, 13th of March, 30th of April, 11th of July, 11th of September, 28th of September, 16th of October, 1st of November and 27th of December.
- We observe that the following events have a decreasing effect on AE demand : 14th of February, 30th of May, 5th of August, 20th of August, 25th of October, 31th of October, 20th of December, 24th of December, 25th of December and 31th of December.
- One of our primary objectives during the development phase was to create a model that could be generalized. Our model can easily be adapted to different settings not only in EDs but also in other health services such as ambulance service or hospitals trying to generate accurate forecast by incorporating special events. Our model could also be generalized to include more special events if necessary.

The proposed model takes i) the historical daily ED attendance and ii) the time series of type of events as inputs and will generate point and probabilis-

tic forecasts as output. We provide the R code for the proposed model and benchmarks used in this paper. They will also be available in an open source R package.

It is important to note that translating forecasts to utility measures is not a straightforward task (Rostami-Tabar et al., 2019). The link between forecast accuracy and staffing in the emergency department needs further investigations to quantify its impact on bottom lines such as calling for external agency, operational costs, staff sickness, patient waiting time and ambulance queue. Moreover, emergency department managers are not familiar with the benefits of probabilistic forecasts. It is crucial to investigate how probabilistic forecasts might be used to inform decisions in the emergency department and how they should be linked to operational planning. We will investigate this in our future work.

References

- Alhusain, F. A. M., Alhassan, N. F. A., Aljohi, W. A. A., Alrumaih, F. I., Al Jerian, N. A., Alharthy, N. A., et al., 2017. Paediatric emergency department during the holidays: Findings from a 10-year analysis of visit rates and trauma patterns. *Journal of Health Specialties* 5 (3), 142.
- Ang, E., Kwasnick, S., Bayati, M., Plambeck, E. L., Aratow, M., 2015. Accurate emergency department wait time prediction. *Manufacturing & Service Operations Management* 18 (1), 141–156.
- Baer, A., Elbert, Y., Burkom, H. S., Holtry, R., Lombardo, J. S., Duchin, J. S., 2011. Usefulness of syndromic data sources for investigating morbidity resulting from a severe weather event. *Disaster medicine and public health preparedness* 5 (1), 37–45.
- Batal, H., Tench, J., McMillan, S., Adams, J., Mehler, P. S., 2001. Predicting patient visits to an urgent care clinic using calendar variables. *Academic Emergency Medicine* 8 (1), 48–53.
- Blunt, I., 2014. Focus on: A&e attendances. *Quality Watch*.

- Booth, A., Harrison, C., Gardener, G., Gray, A., 1992. Waiting times and patient satisfaction in the accident and emergency department. *Emergency Medicine Journal* 9 (2), 162–168.
- Box, G. E., Jenkins, G. M., Reinsel, G. C., Ljung, G. M., 2015. *Time series analysis: forecasting and control*. John Wiley & Sons.
- Boylan, J. E., 2016. Reproducibility. *IMA Journal of Management Mathematics* 27 (2), 107–108.
- Boylan, J. E., Goodwin, P., Mohammadipour, M., Syntetos, A. A., 2015. Reproducibility in forecasting research. *International Journal of Forecasting* 31 (1), 79–90.
- Calegari, R., Fogliatto, F. S., Lucini, F. R., Neyeloff, J., Kuchenbecker, R. S., Schaan, B. D., 2016. Forecasting daily volume and acuity of patients in the emergency department. *Computational and Mathematical Methods in Medicine* 2016.
- Champion, R., Kinsman, L. D., Lee, G. A., Masman, K. A., May, E. A., Mills, T. M., Taylor, M. D., Thomas, P. R., Williams, R. J., 2007. Forecasting emergency department presentations. *Australian Health Review* 31 (1), 83–90.
- DaGar, S., SAHIN, S., YILMAZ, Y., DURAK, U., 2014. Emergency department during long public holidays. *Turkish journal of emergency medicine* 14 (4), 165–171.
- De Livera, A. M., Hyndman, R. J., Snyder, R. D., 2011. Forecasting time series with complex seasonal patterns using exponential smoothing. *Journal of the American Statistical Association* 106 (496), 1513–1527.
- Diehl, A. K., Morris, M. D., Mannis, S. A., 1981. Use of calendar and weather data to predict walk-in attendance. *Southern medical journal* 74 (6), 709–712.
- Durbin, J., Koopman, S. J., 2012. *Time series analysis by state space methods*. Vol. 38. Oxford University Press.

- Faryar, K. A., 2013. The effects of weekday, season, federal holidays, and severe weather conditions on emergency department volume in montgomery county, ohio.
- Gneiting, T., Raftery, A. E., 2007. Strictly proper scoring rules, prediction, and estimation. *Journal of the American Statistical Association* 102 (477), 359–378.
- Hastie, T., Tibshirani, R., Wainwright, M., 2015. *Statistical learning with sparsity: the lasso and generalizations*. CRC press.
- Holleman, D. R., Bowling, R. L., Gathy, C., 1996. Predicting daily visits to a waik-in clinic and emergency department using calendar and weather data. *Journal of General Internal Medicine* 11 (4), 237–239.
- Hong, T., Pinson, P., Fan, S., Zareipour, H., Troccoli, A., Hyndman, R. J., 2016. Probabilistic energy forecasting: Global energy forecasting competition 2014 and beyond. *International Journal of Forecasting* 32, 896–913.
- Hyndman, R. J., Khandakar, Y., 2008. Automatic time series forecasting: the forecast package for R. *Journal of Statistical Software* 26 (3), 1–22.
URL <http://www.jstatsoft.org/article/view/v027i03>
- Hyndman, R. J., Koehler, A. B., 2006. Another look at measures of forecast accuracy. *International journal of forecasting* 22 (4), 679–688.
- Kam, H. J., Sung, J. O., Park, R. W., 2010. Prediction of daily patient numbers for a regional emergency medical center using time series analysis. *Healthcare informatics research* 16 (3), 158–165.
- Kamali, M. F., Tezcan, T., Yildiz, O., 2018. When to use provider triage in emergency departments. *Management Science*.
- Marcilio, I., Hajat, S., Gouveia, N., 2013. Forecasting daily emergency department visits using calendar variables and ambient temperature readings. *Academic emergency medicine* 20 (8), 769–777.

- Reis, B. Y., Mandl, K. D., 2003. Time series modeling for syndromic surveillance. *BMC Medical Informatics and Decision Making* 3 (1), 2.
- Rostami-Tabar, B., Babai, M. Z., Ali, M., Boylan, J. E., 2019. The impact of temporal aggregation on supply chains with arma (1, 1) demand processes. *European Journal of Operational Research* 273 (3), 920–932.
- Rotstein, Z., Wilf-Miron, R., Lavi, B., Shahar, A., Gabbay, U., Noy, S., 1997. The dynamics of patient visits to a public hospital ed: a statistical model. *The American journal of emergency medicine* 15 (6), 596–599.
- Schweigler, L. M., Desmond, J. S., McCarthy, M. L., Bukowski, K. J., Ionides, E. L., Younger, J. G., 2009. Forecasting models of emergency department crowding. *Academic Emergency Medicine* 16 (4), 301–308.
- Sun, Y., Heng, B. H., Seow, Y. T., Seow, E., 2009. Forecasting daily attendances at an emergency department to aid resource planning. *BMC emergency medicine* 9 (1), 1.
- Tai, C.-C., Lee, C.-C., Shih, C.-L., Chen, S.-C., 2007. Effects of ambient temperature on volume, specialty composition and triage levels of emergency department visits. *Emergency Medicine Journal* 24 (9), 641–644.
- Thomas Schneider, A., Luuk Besselink, P., Zonderland, M. E., Boucherie, R. J., Van den Hout, W. B., Kievit, J., Bilars, P., Fogteloo, A. J., Rabelink, T. J., 2018. Allocating emergency beds improves the emergency admission flow. *Interfaces* 48 (4), 384–394.
- Uniejewski, B., Nowotarski, J., Weron, R., 2016. Automated variable selection and shrinkage for day-ahead electricity price forecasting. *Energies* 9 (8), 621.
- Wargon M., Guidet B., H. T. D. H. G., 2009. A systematic review of models for forecasting the number of emergency department visits. *Emergency Medicine Journal* 29, 395–399.
- Xu, K., Chan, C. W., 2016. Using future information to reduce waiting times in the emergency department via diversion. *Manufacturing & Service Operations Management* 18 (3), 314–331.

- Zheng, W., Muscatello, D. J., Chan, A. C., 2007. Deck the halls with rows of trolleys... emergency departments are busiest over the christmas holiday period. *Med J Aust* 187 (11-12), 630–3.
- Ziel, F., Liu, B., 2016. Lasso estimation for gefcom2014 probabilistic electric load forecasting. *International Journal of Forecasting* 32 (3), 1029–1037.
- Ziel, F., Weron, R., 2018. Day-ahead electricity price forecasting with high-dimensional structures: Univariate vs. multivariate modeling frameworks. *Energy Economics* 70, 396–420.